

# Finance-Informed Neural Networks — Deep Learning for Functional Problems in Macroeconomics and Finance

Using Neural Networks to Price Interest Rate Derivatives and for Policy Evaluation in DSGE Models

a thesis written by


Kevin Pierre Mott

submitted in partial fulfillment of the  
requirements for the degree of

**DOCTOR OF PHILOSOPHY**

at the

Tepper School of Business  
Carnegie Mellon University  
Pittsburgh, Pennsylvania  
USA



## Dissertation Committee

**Stephen Spear** (chair)

Tepper School of Business  
Carnegie Mellon University

**Chester Spatt**

Tepper School of Business  
Carnegie Mellon University

**Tetiana Davydiuk**

Carey Business School  
Johns Hopkins University

**Deeksha Gupta** (external reader)

Carey Business School  
Johns Hopkins University

사랑하는 아내에게 바칩니다

## Acknowledgments

This dissertation is a reflection of the many amazing supporters I have in my life.

I first wish to thank Stephen Spear, my advisor and the chair of my dissertation committee. Beyond research advice – of which there was lots – Steve has been involved in the entirety of my doctoral journey. From several courses offered in the first years of my doctoral studies to his generosity supporting conference travel and culminating with his support of my job in Toronto, Steve couldn't have been a better advisor. I am also thankful to rest of my committee: Chester Spatt, Deeksha Gupta, and Tetiana Davydiuk. Any student interested in studying policy with structural methods in a finance department would be lucky to find this committee. Their encouragement of my research agenda, perpetual focus on real-world issues, generosity with their time, and positivity went a long way to keeping my spirits high during the long 6.5 years of graduate school.

It takes a village to raise a PhD student. Aside from my committee, I am grateful for the time and attention of other students and faculty in finance and economics at the Tepper School and beyond. I would like to particularly mention Bryan Routledge for his support as the Area Head of Finance, as well as Burton Hollifield for endless career advice, Bill Hrusa for the opportunity to work with undergraduate math-finance students as a research mentor – an opportunity that helped inspire my third chapter, and Rachel Childers for several courses and many hours of interesting conversations and practical advice. Students particularly crucial to helping navigate my research include Kerry Zhang, Martin Michelini, Xiaonan Hong, and Nick Hoffman.<sup>1</sup> I am also grateful for the research support of my new colleagues at

---

<sup>1</sup>And Ben Iorio, who was always available to power on my desktop.

the University of Toronto whose time I've taken up with some of the work in this dissertation, namely Alexandre Corhay, Jincheng Tong, and David Goldreich. Additionally, Zhigang Feng and Eungsik Kim were great collaborators and mentors throughout my doctoral studies.

Lawrence Rapp and Laila Lee deserve the world for not only putting up with me, but always making my day brighter. Laila, you put a smile on my face every time we talk, whether it's purely shop-talk or (more likely) exchanging stories and photos of our pets. Lawrence, between your sheer administrative prowess and your frank levity, you were always a reliable supporter. Plus, a great unexpected source of Long Island culture when I was homesick.

At even the most stressful periods of graduate school, I was grateful that the experience brought me close to amazing lifelong friends: Zahra Ebrahimi and Martin Michelini, Matt Diabes and Michael Westfall (the four of whom were even in my wedding party), Nick and Geena Hoffman, Dan and Elizabeth Connolly, Jenny Oh, Christine and Beau Dabbs, Kayla Frisoli and Sam Swarts, Avi and Manuela Collis, Sam Levy and Nur Yildirim, Vitaly Mersault, Melda Korkut, Musa Çeldir, Behnam Mohammadi, Daniel de Roux, Mateo Dulce and Laura Muñoz, Orhun and Deniz Gün, Allen Brown, Belle Zhang, Nate Fulham, Jane Pyo, Nilsu Uzunlar, Kerry Zhang, Miguel Oliveira and Maria Aniceto, and Genaro Basulto. The list goes on. For my first three years in Pittsburgh, I was lucky to live with Aaron Satyanarayana, who helped keep me sane. And just next door I was lucky to meet MaryGrace King and Jeremy Beard, filling that role of keeping me sane on Wendover Place after Aaron graduated.

To all of my friends from life before grad school: thanks for sticking around. Research and graduate school can be all-encompassing and isolating. You were all there to provide me with much-needed breaks in the form of phone calls, trips around the world, visiting me often: first in Pittsburgh then in Toronto, and just *being there*. Having friends who understand me so well is something I don't take for granted.

Most of all to my family, without whom I could not have accomplished any of this, I am eternally thankful. I spent a significant amount of my first two years of graduate school at my childhood home due to

COVID-19, and the support of my mom and brother during that time was instrumental in passing my first summer paper and qualifying exams. Since then, we've been lucky enough to expand our family to include Bill, who we all rely on as a bastion of equal parts stability and humor. Even after I moved back to Pittsburgh after the pandemic, my whole family has been there to support me throughout it all. From my dad to all of my grandparents to my aunts and uncles and my cousins, it's hard to overstate how much your support got me through this. I am also lucky enough to have expanded my family during my studies: my in-laws have been understanding and supportive from the first time we spoke on the phone. I love you all.

And last – but certainly not least – to my wife Sae-Seul Park: it's both a cliché and an understatement to say that I couldn't have done this without you. Meeting you during graduate school makes all the years of low salary, research frustration, and uncertainty well worth it. Even when things are objectively hard, you make it all feel so easy. We handled a year of long-distance after you graduated and moved to Toronto without any hiccups. You handled *me* in the last few months before I turned in this very dissertation, which I fear I didn't make easy for you. I feel like I won the lottery in meeting you, and I'm so excited for the rest of our life together.

## Abstract

In all three chapters, I solve functional problems in macroeconomics and finance applications by parameterizing the solution function to the given model with a deep neural network. Following the physics-informed neural networks (PINNs) approach, which embeds partial differential equations directly into the loss function during training, I develop finance-informed neural networks (FINNs). Rather than PDEs governing equilibrium *physical* systems, FINNs incorporate equilibrium conditions from the *financial* model as penalty terms in the loss function. This approach ensures that the trained neural network respects the economic structure of the problem while approximating global policy functions over high-dimensional state spaces without relying on dimensionality reduction or local projection techniques.

### **Student Debt (Forgiveness) in General Equilibrium.**

This chapter evaluates the Biden Administration’s proposed student loan forgiveness policy in a stochastic overlapping generations model with 60 periods of life and three household types differentiated by student debt-to-income ratios. The central finding is that student loan forgiveness generates minimal real economic effects, contradicting the policy’s stated objective of promoting wealth accumulation among over-leveraged borrowers. Despite reduced debt burdens, borrowers allocate the transfer primarily toward consumption rather than retirement savings or productive investment, leaving aggregate capital, production, wages, and asset prices virtually unchanged. For non-borrowers, the policy delivers welfare losses driven almost entirely by higher tax obligations needed to finance the forgiveness, with negligible offsetting gains from general equilibrium spillovers. While forgiveness does provide a small welfare benefit through reduced consumption risk—acting as government-provided intergenerational risk sharing—this effect is quantitatively minor. The results suggest that student loan obligations were not the binding constraint on wealth accumulation for young highly leveraged borrowers, and the forgiveness program operates primarily as a fiscal transfer rather than a mechanism to unlock productive investment.

### **Real and Asset Pricing Effects of Employer Retirement Matching.**

This chapter asks whether employer retirement matching generates meaningful general equilibrium effects on firm investment and output. Employer matching subsidizes household equity purchases, altering savings incentives and potentially changing household intertemporal marginal rates of substitution (MRS). The firm discounts future dividends using the endogenous stochastic discount factor (SDF) arising from household MRS, so matching could affect corporate investment through this repricing chan-

nel. I integrate stochastic overlapping generations and neoclassical q-theory firm investment models, where matching enters households' Euler equations and the household MRS-implied SDF determines the firm's cost of capital. Analytically, I prove in a two-period deterministic model that matching unambiguously increases the SDF, reduces equilibrium returns, and raises capital investment regardless of if the matching is financed out of the labor or capital share of output: households tolerate lower market returns because their effective returns inclusive of the match remain attractive. Solving the full 60-period stochastic model using FINNs confirms these predictions quantitatively. Introducing empirically realistic matching reduces equilibrium equity returns by 79 basis points, increases the aggregate capital stock by 6.1%, and raises wages by 1.7%.

### **Deep Learning the Term Structure for Derivatives Pricing.**

This chapter introduces a no-arbitrage, Monte Carlo-free approach to pricing path-dependent interest rate derivatives. The Heath-Jarrow-Morton model gives arbitrage-free contingent claims prices but is infinite-dimensional, making traditional numerical methods computationally prohibitive. To make the problem computationally tractable, I cast the stochastic pricing problem as a deterministic partial differential equation (PDE). Finance-Informed Neural Networks (FINNs) solve this PDE directly by minimizing violations of the differential equation and boundary condition, with automatic differentiation efficiently computing the exact derivatives needed to evaluate PDE terms. FINNs achieve pricing accuracy within 0.04 to 0.07 cents per dollar of contract value compared to Monte Carlo benchmarks. Once trained, FINNs price caplets in a few microseconds regardless of dimension, delivering speedups ranging from 300,000 to 4.5 million times faster than Monte Carlo simulation as the state space discretization of the forward curve grows from 10 to 150 nodes. The major Greeks—theta and curve deltas—come for free, computed automatically during PDE evaluation at zero marginal cost, whereas Monte Carlo requires complete re-simulation for each sensitivity. The framework generalizes naturally beyond caplets to other path-dependent derivatives—caps, swaptions, callable bonds—requiring only boundary condition modifications while retaining the same core PDE structure.

# Contents

|                                                                          |     |
|--------------------------------------------------------------------------|-----|
| Dissertation Committee                                                   | i   |
| Acknowledgments                                                          | iii |
| Abstract                                                                 | vi  |
| 1 Student Debt (Forgiveness) in General Equilibrium                      | I   |
| 1.1 Introduction                                                         | I   |
| 1.1.1 Institutional Setting of Student Loans                             | 5   |
| 1.1.2 Model Overview                                                     | 6   |
| 1.1.3 Preview of Results                                                 | 8   |
| 1.2 Related Literature                                                   | 9   |
| 1.2.1 Higher Education Financing                                         | 9   |
| 1.2.2 Overlapping Generations Models, Policy Analysis, and Asset Pricing | 10  |
| 1.2.3 Computational Methods                                              | 11  |
| 1.3 The Model                                                            | 12  |
| 1.3.1 Market Structure                                                   | 13  |
| 1.3.2 Households                                                         | 14  |
| 1.3.3 Firms                                                              | 16  |
| 1.3.4 The Government                                                     | 17  |
| 1.3.5 Market Clearing                                                    | 17  |
| 1.3.6 Equilibrium                                                        | 17  |
| 1.4 Data Sources and Calibration                                         | 19  |
| 1.4.1 Data Source                                                        | 19  |
| 1.4.2 Estimating Student Debt Balances After Forgiveness                 | 21  |
| 1.5 Solution Method: FINNs                                               | 25  |
| 1.5.1 Applications of Domain Knowledge                                   | 25  |
| 1.5.2 FINN Architecture and Training                                     | 28  |
| 1.6 Results                                                              | 32  |
| 1.7 Conclusion                                                           | 38  |
| 2 Real and Asset Pricing Effects of Employer Retirement Matching         | 41  |
| 2.1 Introduction                                                         | 41  |
| 2.1.1 Institutional Setting of Firm Retirement Contributions             | 43  |
| 2.1.2 Preview of Results                                                 | 44  |
| 2.2 Related Literature                                                   | 47  |

|       |                                                                         |     |
|-------|-------------------------------------------------------------------------|-----|
| 2.2.1 | Overlapping Generations Models and Retirement Policy . . . . .          | 47  |
| 2.2.2 | Household Portfolio Choice and Retirement Savings Behavior . . . . .    | 48  |
| 2.2.3 | Asset Pricing in Overlapping Generations Economies . . . . .            | 49  |
| 2.2.4 | Neoclassical Firm Investment and the Cost of Capital . . . . .          | 50  |
| 2.2.5 | Computational Methods: Deep Learning for Economics . . . . .            | 51  |
| 2.3   | The Model . . . . .                                                     | 52  |
| 2.3.1 | Market Structure . . . . .                                              | 52  |
| 2.3.2 | Households . . . . .                                                    | 53  |
| 2.3.3 | Firms . . . . .                                                         | 55  |
| 2.3.4 | Market Clearing . . . . .                                               | 56  |
| 2.3.5 | Equilibrium . . . . .                                                   | 56  |
| 2.4   | Two-Period Steady State Equilibrium Analysis . . . . .                  | 58  |
| 2.4.1 | Households . . . . .                                                    | 58  |
| 2.4.2 | Representative Firm . . . . .                                           | 60  |
| 2.4.3 | Marginal vs. Average Tobin's Q . . . . .                                | 62  |
| 2.4.4 | Worker-Financed Match . . . . .                                         | 63  |
| 2.5   | Sixty-Period Computation . . . . .                                      | 65  |
| 2.5.1 | Data Sources and Calibration . . . . .                                  | 66  |
| 2.5.2 | Solution Method: FINNs . . . . .                                        | 68  |
| 2.5.3 | Applications of Domain Knowledge . . . . .                              | 69  |
| 2.5.4 | FINN Architecture and Training . . . . .                                | 72  |
| 2.5.5 | Computation of an Efficient SDF for the Firm's Euler Equation . . . . . | 77  |
| 2.5.6 | Results . . . . .                                                       | 79  |
| 2.6   | Conclusion . . . . .                                                    | 84  |
| 3     | Deep Learning the Term Structure for Derivatives Pricing . . . . .      | 87  |
| 3.1   | Introduction . . . . .                                                  | 87  |
| 3.2   | Literature Review . . . . .                                             | 90  |
| 3.3   | Heath-Jarrow-Morton Model Refresher . . . . .                           | 92  |
| 3.3.1 | Musiela Parameterization . . . . .                                      | 94  |
| 3.4   | Standard Procedure for Computing the HJM Model . . . . .                | 95  |
| 3.4.1 | Data and Forward Curve Construction . . . . .                           | 95  |
| 3.4.2 | Volatility Estimation . . . . .                                         | 96  |
| 3.4.3 | Computing the No-Arbitrage Drift . . . . .                              | 97  |
| 3.4.4 | Local Volatility Specification . . . . .                                | 98  |
| 3.5   | Feynman-Kac Formula in the HJM Framework . . . . .                      | 100 |
| 3.6   | Application: Caplet Pricing . . . . .                                   | 104 |
| 3.6.1 | Neural Network Strategy . . . . .                                       | 105 |
| 3.6.2 | FINN Architecture and Training . . . . .                                | 107 |
| 3.6.3 | Results . . . . .                                                       | 112 |
| 3.7   | Conclusion . . . . .                                                    | 119 |
|       | References . . . . .                                                    | 121 |

# Chapter 1: Student Debt (Forgiveness) in General Equilibrium

## 1.1 Introduction

Student debt has grown to a magnitude of approximately 1.8 trillion dollars, the majority of which is issued by the United States federal government. At the household level, student debt balances constitute the largest liability aside from mortgages. While there is widespread agreement that higher education can provide a pathway to higher earnings, the necessity of graduates to finance education with debt has led to student debt balances and earnings being positively correlated in the data (Catherine and Yannelis 2023). Obscured by that positive correlation is the fact that there exist low-earning, high-student-debt households, burdened by substantial balances and facing little prospect of relief under the stringent repayment clauses governing federal student loans. For example, delinquent borrowers face automatic wage garnishment and tax rebate sequestration, and the debt is quite rarely discharged even in bankruptcy proceedings.

Student loan forgiveness policies raise important questions for macroeconomic research. At the individual level, the credit constraints faced by young would-be borrowers provide justification for government intervention in the form of student loan subsidization. But in the aggregate, does a productivity externality of education justify this government subsidization financed indirectly by non-borrowers? Along the same vein, does this productivity externality provide a general equilibrium channel by which even *non-*

*borrowers* may benefit from a loan forgiveness program, by freeing highly-educated borrowers previously saddled with debt to engage in more productive financial and real activities? The intertemporal tension between student loans when young and retirement portfolios when old is an important aspect of the economic life-cycle in the contemporary American education and financial systems. As mentioned in policy debates, the intention of forgiveness is to alleviate financial pressures on middle-class borrowers for the sake of building wealth through buying homes, saving for retirement, and starting small businesses. This paper focuses on the channel of retirement savings and capital accumulation. In particular, I quantify the general equilibrium effects in asset markets that may either amplify or attenuate the redistributive nature of student loan forgiveness.

The policy environment examined in this paper is the Biden Administration's August 24, 2022 announcement of up to \$10,000 in debt cancellation for non-Pell Grant recipients and up to \$20,000 for Pell Grant recipients, subject to income thresholds of \$125,000 (single) or \$250,000 (married couples). According to the White House press release:

Middle-class borrowers struggle with high monthly payments and ballooning balances that make it harder for them to build wealth, like buying homes, putting away money for retirement, and starting small businesses.

Subsequent legal challenges culminated in a June 30, 2023 Supreme Court decision (*BIDEN v. NEBRASKA*, No. 22–506) blocking the program. Nevertheless, the Biden Administration proceeded with piecemeal forgiveness for certain groups, approving nearly \$138 billion in cancellation for almost 3.9 million borrowers through more than two dozen executive actions as of this writing. Despite the Supreme Court ruling, I choose to model the ramifications of the proposed policy in a dynamic stochastic general equilibrium framework to understand its economic mechanisms and distributional consequences.

By including households who do not hold student debt, I can quantify the effect that this policy has on non-borrowers through general equilibrium channels. Furthermore, as household indebtedness has been rising across multiple credit markets (credit cards, mortgages, auto loans), I provide a flexible frame-

work with which to consider the general equilibrium effects of debt relief policies more broadly. The central questions I seek to address are: Does student loan forgiveness enable borrowers to accumulate wealth faster through increased retirement savings or productive investment? Do non-borrowers benefit or suffer from such a policy through general equilibrium effects on wages, asset prices, and returns? What role do asset pricing channels play in redistributing the policy's effects across generations and household types?

To answer these questions, I develop a fully stochastic overlapping generations model with 60 periods of life and three household types differentiated by their student debt-to-income ratios. The model features aggregate productivity risk, endogenous capital accumulation, and government-issued student loans financed through a balanced budget. This framework is essential for studying policies that transfer resources across age cohorts and may alter the characteristics of the marginal investor in asset markets. As argued by Glover et al. (2020), young households have little financial wealth compared to human wealth (the present-value of future labor earnings), while the opposite is true for older households. Student loan balances are disproportionately held by younger households. These facts imply that a student loan forgiveness program amounts to a transfer of financial wealth to young households, thereby shifting the characteristics of the marginal investor. Closely following the argument of Constantinides, Donaldson, and Mehra (2002), the transfer of financial wealth to young, financially constrained households can substantially affect the stochastic discount factor and lead to redistribution through the channel of heterogeneous portfolios of financial wealth.

Despite the growing macroeconomic and household-level importance of student debt, there is a lack of structural general equilibrium models which incorporate student debt as a financial liability, so the *asset pricing* effects of student debt policies remain poorly understood. In my model, I take as given the educational choices and corresponding student debt levels of the young, reflecting the reality that partial loan forgiveness policies target those who have *already* made their education-loan decisions and are now in the workforce. This backwards-looking focus allows me to isolate the redistributive effects of the policy on current borrowers without conflating them with forward-looking changes to educational

investment or borrowing behavior. A study of policies that affect endogenous education choice and the associated moral hazard problems—whereby future borrowers might anticipate forgiveness and alter their borrowing decisions, or universities might raise tuition in anticipation of debt relief—is left for future work.

I find that student loan forgiveness fails to achieve its stated objective of promoting wealth accumulation among borrowers. Despite the reduction in debt burdens, borrowers allocate the transfer primarily toward increased consumption rather than retirement savings or productive investment. Consequently, aggregate capital, production, wages, and asset prices remain virtually unchanged across policy environments. For non-borrowers, the policy generates welfare losses driven almost entirely by higher tax obligations needed to finance the forgiveness, with negligible offsetting benefits from general equilibrium effects. In a counterfactual policy experiment wherein the tax obligation of non-borrowers is pinned at pre-forgiveness levels and the tax bill is solely borne by those who borrowed student loans, the welfare levels of non-borrowers are almost unchanged and positive welfare gains are observed for both low- and high-loan-to-income borrowers. While forgiveness does provide a small welfare gain through reduced consumption risk for all household types—acting as a form of government-provided intergenerational risk sharing—this effect is quantitatively minor and does not offset the fiscal costs borne by non-borrowers.

These results suggest that student loan obligations were not the binding constraint on wealth accumulation for young borrowers. The borrowers with high loan-to-income ratios are not freed up to invest productively by the forgiveness, instead preferring to consume almost all of the transfer, leaving aggregate capital and production largely unchanged. To this extent, the stated goal of the policy to enable faster wealth accumulation of borrowers is not successful. The findings underscore that the Biden Administration’s forgiveness proposal functions primarily as a fiscal transfer rather than a mechanism to unlock productive investment or alter macroeconomic outcomes.

### 1.1.1 Institutional Setting of Student Loans

For context as to the magnitude of student loan balances, it is worth noting that approximately 89% of student loan balances outstanding are owed to the federal government throughout the 2010s and 2020s. The total balance of student loans owned and securitized is reported for the first quarter of 2022 at \$1.75 trillion, which compares to \$40.72 trillion total market capitalization of U.S. equities. Total student debt balances outstanding in the United States have increased when measured in level, as a percentage of GDP, and as a portion of total household debt, and compared to total US equity markets capitalization. At the household level, average federal loan borrowing amounts to \$24,280 for undergraduates and \$45,680 for graduates as of the 2020-2021 academic year (Ma and Pender 2021; Board of Governors of the Federal Reserve System [US] 2022). Federal student loans are federally issued at a fixed rate of interest set by the US Congress (Ma and Pender 2021).

A distinctive feature of the federal student loan system is the entitlement-like nature of loan issuance. Federally-issued student loans are granted to anyone who qualifies based on higher-education enrollment, without regard to field of study, expected post-graduation earnings, credit history, or ability to repay. This stands in sharp contrast to private credit markets, where lenders condition loan terms on borrower risk characteristics and expected income streams. The absence of underwriting standards and the unlimited availability of federal loans up to statutory limits create a unique environment in which students can accumulate large debt positions irrespective of their future earnings potential. A prospective philosophy major faces the same borrowing terms as a prospective engineering major, despite vastly different expected earnings trajectories. Similarly, students enrolling in institutions with poor labor market outcomes for graduates face no additional scrutiny or pricing adjustments relative to students at elite institutions with strong placement records. This entitlement structure, combined with statutory limits on borrowing that have increased substantially over time, enables debt accumulation patterns that may be inappropriate for the forecasted lifetime earnings of some borrowers. The resulting disconnect between debt obligations and earnings capacity is a defining characteristic of the student debt distribution and motivates the policy concerns that this paper addresses. Indeed, as documented in the Survey of Consumer Finances and

other household surveys, the data reveal a substantial population of high-debt, low-income households—borrowers with large outstanding balances relative to their realized earnings—who are likely the primary intended beneficiaries of the Biden Administration’s forgiveness proposal.

### 1.1.2 Model Overview

Since student debt is awarded to the young and the central tension I seek to analyze is the tradeoff between student debt financing and retirement savings, I use the overlapping generations (OLG) framework pioneered by Samuelson (1958). The model features three household types per generation: non-borrowers and two types of borrowers differentiated by their loan-to-income ratios. This classification is motivated by empirical patterns in the Survey of Consumer Finances and related household surveys: the highest-leveraged households in student debt—those with the largest debt-to-income ratios—systematically exhibit lower realized earnings and face worse economic prospects compared to borrowers with similar absolute debt levels but higher incomes. By segmenting households along the loan-to-income dimension, I capture the financially distressed population most likely targeted by forgiveness policies—those for whom debt service obligations represent a large share of disposable income and whose debt burdens are most misaligned with earnings capacity. Each type within each generation is populated by a continuum of identical agents, allowing me to refer to the entire generation as ‘the (representative) agent (for that generation of that type).’ The life-cycle profile of income as estimated from common data sources implies a low-income retirement period; by choosing standard utility functions, younger generations will find it optimal to save for retirement financing, guaranteeing non-zero asset demand schedules.

Households accumulate productive capital over their lives, renting capital stocks to the representative firm for production. Student debt forgiveness frees up financial resources in indebted households, potentially allowing for higher capital accumulation—which can be thought of as a reduced-form proxy for entrepreneurship activities or other productive investment opportunities undertaken through unmodeled financial intermediaries. I assume that agents are financially constrained, such that student loans and private debt are not redundant financial liabilities. Since my interest lies with the policy’s effects on cur-

rent graduates, I endow agents with human capital-student debt tuples calibrated to the data on life-cycle debt-to-income profiles, abstracting from the effects of this policy on future educational investment or borrowing decisions.

The OLG setting is a natural framework in which to explore the interaction between student loans and financial assets. While social security—richly studied in the OLG literature—is a program of transfers to the oldest in society, subsidized student loan issuance and forgiveness are transfers to the youngest. Furthermore, the history dependence on endogenous variables in OLG models generates intergenerational heterogeneity even in the absence of idiosyncratic shocks, through the inability of agents to engage in trade with already dead or yet-unborn agents—a phenomenon termed ‘restricted market participation’ (Cass and Shell 1983). Additionally, the realization of an aggregate shock in each period induces agents to rebalance their portfolios, introducing history-dependent shocks to rational expectations equilibrium asset returns (Spear and Srivastava 1986; Duffie et al. 1994; Citanna and Siconolfi 2010). This generational risk is (privately) uninsurable, opening an opportunity for a social planner to improve on private market outcomes.

Consumption-based asset pricing is therefore a natural framework with which to ascertain the general equilibrium asset pricing effects of student loan policies. Since the decisions of any agent in a general equilibrium environment affect the outcome of all agents, student loans ‘matter’ in some sense even to non-borrowers. But I posit that the link between borrowers and non-borrowers is more direct: a more highly educated population is perhaps more productive, so that even non-college attendees would benefit from higher education levels. It is in this spirit that a general equilibrium exploration of the *indirect* redistributive effects of student loan policies is in order. I identify several plausible channels through which student debt forgiveness may affect lifetime utility for both borrowers and non-borrowers: direct wealth effects for borrowers, retirement portfolio choice decisions for all agents, general equilibrium wage effects through aggregate capital accumulation, and productivity externalities of educated households making capital investment decisions. The model allows me to quantify the relative importance of these channels and assess whether the indirect redistributive effects through general equilibrium prices are

quantitatively significant.

### 1.1.3 Preview of Results

I find that student loan forgiveness fails to achieve its stated objective of promoting wealth accumulation among borrowers. Despite the reduction in debt burdens, borrowers allocate the transfer primarily toward increased consumption rather than retirement savings or productive investment. This behavioral response is notable given the aggressive calibration of precautionary savings motives in the model (coefficient of relative risk aversion  $\gamma = 3$ ), suggesting that life-cycle consumption smoothing incentives dominate precautionary savings incentives. Consequently, aggregate capital, production, wages, and asset prices remain virtually unchanged across policy environments. The minimal effects on aggregate quantities imply that the channels through which student debt forgiveness might indirectly benefit non-borrowers—higher wages from increased aggregate capital, altered asset pricing from shifts in the marginal investor, or productivity externalities from deleveraged entrepreneurial borrowers—fail to materialize in quantitatively meaningful ways.

For non-borrowers, the policy generates welfare losses driven almost entirely by higher tax obligations needed to finance the forgiveness, with negligible offsetting benefits from general equilibrium effects. A counterfactual experiment holding non-borrower taxes fixed at baseline levels confirms this interpretation: when non-borrowers are insulated from the fiscal burden, their welfare remains essentially unchanged, indicating that general equilibrium price effects are quantitatively unimportant. While forgiveness does provide a small welfare gain through reduced consumption risk for all household types—acting as a form of government-provided intergenerational risk sharing by stabilizing borrower consumption responses to aggregate shocks—this effect is quantitatively minor and does not offset the fiscal costs borne by non-borrowers. Overall, the results suggest that student loan obligations were not the binding constraint on wealth accumulation for young borrowers, and the forgiveness program operates primarily as a fiscal transfer rather than a mechanism to unlock productive economic activity.

## 1.2 Related Literature

This paper builds on three interconnected strands of literature: empirical and structural research on student debt and higher education finance, overlapping generations models with aggregate risk and asset pricing, and computational methods for solving high-dimensional dynamic equilibrium models.

### 1.2.1 Higher Education Financing

The literature on student debt in finance and economics has been growing rapidly. Early structural work by Ionescu (2009) develops a model of student borrowing with default risk. Establishing that young borrowers are credit constrained, Lochner and Monge-Naranjo (2011a) and Lochner and Monge-Naranjo (2011b) provide evidence that student loans are a worthwhile borrowing vehicle and examine how credit constraints shape human capital investment decisions. More recent contributions study the returns to education across different fields: Bleemer and Mehta (2022) document substantial heterogeneity in earnings by major, reinforcing the concern that the entitlement structure of federal loans may enable debt accumulation patterns misaligned with earnings capacity. Abbott et al. (2019) provide evidence on the education-to-debt dynamics over the life cycle. To unpack the distributional effects of the Biden Administration’s forgiveness policy specifically, Catherine and Yannelis (2023) use micro-level survey data to compute novel measures of the net present value of forgiveness among borrowers. Without an equilibrium mechanism in mind, however, they are unable to quantify the redistributive effects the policy has on non-borrowers through general equilibrium channels—a central focus of my paper. Related empirical work includes Chakrabarti et al. (2020), which shows a negative relationship between student debt balances and performance in higher education, and Morazzoni (2022), which studies the link between student debt overhang and entrepreneurship. The survey of the empirical student debt literature from Yannelis and Tracey (2022) is more thorough than my own. Most closely related to my structural approach, Fu, Lin, and Tanaka (2025) build a model exploring the tradeoff between costly human capital investment (on-the-job learning) and loan repayment, finding that income-based repayments increase both welfare and government revenue despite losses to the loan program revenue itself. While my pa-

per evaluates the one-time forgiveness proposal, Boutros, Clara, and Gomes (2024) study optimal student loan contract design by comparing alternative payment schedules that offer partial or full deferral of repayment until later in life when marginal utility is lower, finding large welfare gains comparable to debt relief but without adverse fiscal implications. My paper complements this literature by explicitly modeling general equilibrium asset pricing effects and the role of non-borrowers in a stochastic OLG framework.

### 1.2.2 Overlapping Generations Models, Policy Analysis, and Asset Pricing

The use of overlapping generations models to study fiscal policy and asset pricing has a long tradition in macroeconomics. The classic contributions of Auerbach and Kotlikoff (1987) established the OLG framework as a powerful tool for evaluating intergenerational redistribution from tax and transfer policies. In the asset pricing literature, Storesletten, Telmer, and Yaron (2007) explore the extent to which idiosyncratic shocks over the life cycle matter to asset prices, while Hasanhodzic (2015) show that soft borrowing costs, rather than hard borrowing constraints, can generate realistic risk premia in OLG economies. Glover et al. (2020) use the Great Recession as a setting to explore how asset prices contribute to redistribution across the life cycle, and include a good survey of OLG-based asset pricing. The mechanism linking the age distribution of wealth holders to equilibrium asset returns has been studied by Gârleanu and Panageas (2015) under the title “young, old, conservative, and bold,” and by Geanakoplos, Magill, and Quinzii (2004), who document how demography drives long-run predictability of the stock market. My paper contributes to this literature by studying how student loan forgiveness—a transfer to the young—affects the identity of the marginal investor and thereby equilibrium asset prices, following the theoretical insights of Constantinides, Donaldson, and Mehra (2002) on how transfers of financial wealth to financially constrained young households can substantially affect the stochastic discount factor. Works I saw early in graduate school that were foundational in generating my research agenda include Moretto (2021) and Kim (2018).

### 1.2.3 Computational Methods

I refer the reader to Fernández-Villaverde (2025) for the most up-to-date comprehensive survey of deep learning methods in macroeconomics. In the computational macroeconomics literature, my methodology most closely resembles that of Azinovic, Gaegauf, and Scheidegger (2019) in my use of deep neural networks to approximate global policy functions. The use of neural networks allows me to use the complete state variable without truncation, due to the cheap evaluation of neural networks. This is in contrast to the model-reduction literature pioneered by Krusell and Smith (1998), which seeks a sufficient statistic to represent the state variable, in practice boiling down to a finite set of moments regarding the distribution that comprises the state variable. Related deep learning approaches include L. Maliar, S. Maliar, and Winant (2021) and Han, Yang, and E (2021), the latter of whom use deep learning to endogenously learn the most economically informative moments in a model reduction approach. In continuous time financial intermediate macro-finance models, cutting-edge methods are in the vein of Gopalakrishna (2021), who incorporates active learning to endogenously oversample regions that are hard for the machine learning algorithm to learn. By embedding economic equilibrium conditions directly into the loss function—an approach inspired by physics-informed neural networks (Raissi, Perdikaris, and Karniadakis 2019)—I am able to solve for equilibrium in a high-dimensional stochastic OLG economy with heterogeneous agents and aggregate shocks without resorting to approximations of the state space.

### 1.3 The Model

I use an overlapping generations model with  $J = 3$  types per generation.<sup>1</sup> Indexed by  $j \in \{0, \dots, J-1\}$ , I consider type  $j = 0$  to be households with no student debt, type  $j = 1$  to be households in the bottom half of the loan-to-income distribution (conditional on holding student debt) and type  $j = 2$  to be households in the top half of the loan-to-income distribution. This type classification along the loan-to-income dimension is motivated by the same empirical patterns discussed in the Model Overview: the highest-leveraged households in student debt—those with the largest debt-to-income ratios—systematically exhibit lower realized earnings and face worse economic prospects compared to borrowers with similar absolute debt levels but higher incomes. By segmenting households in this way, I capture the financially distressed population most likely targeted by the Biden Administration’s forgiveness policy—those for whom debt service obligations represent a large share of disposable income and whose debt burdens are most misaligned with earnings capacity. This classification allows me to isolate the effects of forgiveness on the borrowers who are most financially constrained by their student loan obligations and most in need of relief according to the policy’s stated goals.

I model the life cycle from entrance to the workforce through retirement, so that the education choice (and therefore loan balances and human capital stocks) are endowed, set exogenously to match the data. A model with microfounded education, loan, and human capital choices is left to later research. Student debt payments and balances will be calibrated by the data. By taking the payments and balances as governed by the data, I can abstract from heterogeneity in payment systems in the data, such as enrollment in various payment programs offered by the government and differential interest rate schedules faced by different borrowers. This substantially reduces the set of assumptions required to model the large variety of repayment programs offered to borrowers.

There are  $I$  periods of life for each generation. Lifetimes are deterministic.

Available for purchase at all stages of life are financial assets for the purpose of consumption smooth-

---

<sup>1</sup>In this and other indexing choices, I start my indices at 0.

ing and retirement financing, taking the form of productive capital. Discussion of these assets is also included below. Households rent their stock of capital to firms and inelastically supply labor in exchange for competitively determined rental returns on capital and competitively determined wages for labor. The financing of productive capital provides a plausible channel for a positive externality of debt-financed education: by borrowing to pay for college, households can begin accumulating capital earlier and thereby growing the aggregate capital stock and the output of the economy. This is precisely one of the general equilibrium forces that will be studied in this paper.

Since I model student debt load as an exogenous endowment, student debt issuance to the youngest households of each type in any period are financed by the government budget constraint: it will be paid for by tax revenue and repayment on previously issued loans.

For any variable  $x$ , subscripts  $x_{i,j,t}$  means the value of  $x$  for the  $i$ -th oldest generation of type  $j$ ; i.e. the  $j$ -type generation born at  $t - i$ . Additionally, for any lower-case  $i, j, t$ -subscripted variable, let the upper-case script variable denote its  $ij$ -vector:  $\mathcal{X}_t := \{x_{i,j,t}\}_{i,j}$ .

### 1.3.1 Market Structure

The source of aggregate risk in the economy will be the total factor productivity affecting firm output, an autoregressive process of degree one –  $AR(1)$  – in logs. It is well known that models of this type have a Markov equilibrium with history dependence in the state variable (Spear 1988). As shown by Henriksen and Spear (2012) multi-period-lived OLG models with infinitely-lived financial assets subject to aggregate uncertainty generically lack strongly stationary (or, in the terminology of Citanna and Siconolfi (2007), short-memory) equilibria, thus producing *endogenously incomplete markets*.

Log-TFP follows the following process:

$$\log Z_{t+1} = \rho \log Z_t + \sigma_\epsilon \epsilon_{t+1} \quad (1.1)$$

where  $0 < \rho < 1$  so that  $\log Z$  has a stationary distribution and the shock  $\epsilon$  is drawn from a standard

Normal Distribution:  $\epsilon \sim \mathcal{N}(0, 1)$ .

As discussed above, I model one asset. Capital is accumulated by households, depreciating at constant rate  $\delta$  per period, and rented to the firm at competitive rental rates  $r_t$  in period  $t$ . Agent of age  $i$  and type  $j$  at time  $t$  chooses how much capital to accumulate for the next period:  $k_{i+1,j,t+1}$ . The net return on capital is given as  $R_t = 1 + r_t - \delta$ .

Short sales on capital are forbidden, so that  $k_{i,j,t} \geq 0$  for all  $i, j, t$ .

### 1.3.2 Households

As mentioned above, each generation lives for exactly  $I$  periods. Using  $t$  as the time index, the index  $i \in \{0, \dots, I - 1\}$  represents age concurrently with time. The index  $j \in \{0, \dots, J - 1\}$  represents types.

Agents are endowed with labor efficiency units  $\ell_{i,j}$  over their life-cycle for all  $i, j, t$ . Types  $j$  are also endowed with initial student debt balance  $d_{0,j}$  to finance education. The debt  $d_{i,j}$  is paid back in payments  $a_{i,j}$  by generation  $i$  of type  $j$  in period  $t$  over the life-cycle.

In each period, agents choose optimal consumption and savings in risky capital, subject to the constraints as described in the previous section.

Households earn income from wages:

$$y_{i,j,t} = w_t \ell_{i,j} \tag{1.2}$$

at wage rate  $w_t$  competitively determined per efficiency unit of labor.

The government institutes a linear tax at rate  $\tau_t$  to finance its spending on new student debt issuance, more discussion of which will be included below in the **Government** section.

The sequential budget constraints for agents alive at time  $t$  are thus:

$$c_{i,j,t} + i_{i,j,t} = y_{i,j,t}(1 - \tau_t) + r_t k_{i-1,j,t-1} - a_{i,j,t} \quad (1.3)$$

$$i_{i,j,t} = k_{i+1,j,t+1} - (1 - \delta)k_t \quad (1.4)$$

Since agents die with certainty after their last period of life and are endowed with no financial assets:

$$k_{0,j} = k_{I-1,j} = 0 \text{ for all } j, t.$$

Households form utility over lifetime consumption:

$$U : \mathbb{R}_{++}^I \rightarrow \mathbb{R}$$

I'll assume additively time-separable von Neumann–Morgenstern expected constant relative risk aversion

(CRRA) utility, with coefficient of relative risk aversion  $\gamma$ :

$$U(\{c_{i,j,t+i}\}_{i=0}^{I-1}) = \mathbb{E}_t \sum_{i=0}^{I-1} \beta^i u(c_{i,j,t+i}) : u(c) = \begin{cases} \frac{c^{1-\gamma}-1}{1-\gamma} & \gamma \in \mathbb{R}_+ \setminus \{1\} \\ \ln(c) & \gamma = 1 \end{cases}$$

Then the households' formal problem can be written as:

$$\max_{\{c_{i,j,t+i}, k_{i+1,j,t+i+1}\}_{i=0}^{I-2}} \left\{ U(\{c_{i,j,t+i}\}_{i=0}^{I-1}) \right\} \text{ s.t.} \quad (1.5)$$

$$c_{i,j,t} + i_{i,j,t} = y_{i,j,t}(1 - \tau_t) + r_t k_{i-1,j,t-1} - a_{i,j,t} \quad (1.6)$$

$$i_{i,j,t} = k_{i+1,j,t+1} - (1 - \delta)k_t \quad (1.7)$$

$$k_{i+1,j,t+1} \geq 0 \quad (1.8)$$

Let  $\mu_{i,j,t}^k$  be the Lagrange multiplier on the capital nonnegativity constraint for household  $i, j$  at period  $t$ . Households' optimality conditions are given by the following Euler equations and KKT condi-

tions:

$$u'(c_{i,j,t}) = \beta \mathbb{E}_t [u'(c_{i+1,j,t+1}) (1 + r_{t+1} - \delta)] + \mu_{i,j,t}^k \quad (1.9)$$

$$k_{i+1,j,t+1} \geq 0 \quad (1.10)$$

$$\mu_{i,j,t}^k k_{i+1,j,t+1} = 0 \quad (1.11)$$

### 1.3.3 Firms

The economy is populated by a representative firm. The firm hires all types of workers. Firms are endowed with period production function  $F(\cdot, \cdot) : \mathbb{R}_+^2 \rightarrow \mathbb{R}_+$ , which has two components: capital  $K_t$  and labor  $L_t$ .

The firm is subject to the aggregate productivity shock  $Z_t$  and has Cobb-Douglas technology with parameter  $\alpha$  so that output  $Y_t$  is given by:

$$Y_t = Z_t F(K_t, L_t) = Z_t \bar{Z} K_t^\alpha L_t^{1-\alpha} \quad (1.12)$$

The firm is perfectly competitive, choosing capital stock and labor demand and then paying competitive wages to clear the market. Therefore the firm's problem is given by:

$$\max_{K_t, L_t} \{ Z_t \bar{Z} K_t^\alpha L_t^{1-\alpha} - w_t L_t - r_t K_t \}$$

with optimality conditions to clear the markets for labor and capital:

$$w_t = (1 - \alpha) Z_t \bar{Z} \left( \frac{K_t}{L_t} \right)^\alpha \implies L_t = \sum_{i,j} \ell_{i,j,t} \quad (1.13)$$

$$r_t = \alpha Z_t \bar{Z} \left( \frac{L_t}{K_t} \right)^{1-\alpha} \implies K_t = \sum_{i,j} k_{i,j,t} \quad (1.14)$$

### 1.3.4 The Government

The government collects income tax revenue from the linear tax schedule  $T_t = \sum_{i,j} \tau_t y_{i,j,t}$  and student debt repayments  $A_t = \sum_{i,j} a_{i,j,t}$ . The government issues student debt balances  $D_t = d_{0,j}$  to the youngest households of each type.

Therefore the government's budget constraint is given as:

$$D_t = A_t + T_t \quad (1.15)$$

This can be re-written more conveniently to immediately solve for the tax level:

$$\tau_t = \frac{D_t - A_t}{\sum_{i,j} y_{i,j,t}} \quad (1.16)$$

### 1.3.5 Market Clearing

The capital and labor markets need to clear:

$$\sum_{i,j} k_{i,j,t} = K_t \quad (1.17)$$

$$\sum_{i,j} \ell_{i,j} = L_t \quad (1.18)$$

The consumption market will clear via Walras's Law.

### 1.3.6 Equilibrium

The equilibrium of study in long-lived stochastic overlapping generations models is that of a recursive Markov equilibrium.

**Definition 1.** *The **recursive Markov equilibrium** is defined by time-homogeneous policy functions for*

*the purchase and pricing of assets for state variable  $\chi$ :*

$$\{K'(\chi), \mu^k(\chi)\}$$

*where  $\chi \supseteq (K, Z)$  is taken to be at a minimum the lagged asset holdings of all agents and the current realization of the shock. The time-homogeneous policy functions solve the households' problems, the firm's problem, and the government budget constraint. Markets for capital, bonds, and labor must clear. Feasibility arises from Walras's Law and is given by*

$$Y_t = C_t + D_t + I_t$$

*where  $C_t$  is aggregate consumption,  $I_t$  is aggregate investment, and  $D_t$  is issuance of new student loans.*

## 1.4 Data Sources and Calibration

The main model objects that need external calibration from the data are the life-cycle profiles of labor efficiency units, student debt balances, and student debt payments. Further details on those below. Other model parameters are calibrated as follows:

Table 1.1: Calibrated Parameters

| Description                  | Symbol     | Value            | Target / Source           |
|------------------------------|------------|------------------|---------------------------|
| Periods of life              | $I$        | 60               | –                         |
| Efficiency units of labor    | $\ell_i$   | See Figure (1.1) | SCF (2019)                |
| Discount factor              | $\beta$    | 0.950            | –                         |
| Relative risk aversion       | $\gamma$   | 3                | –                         |
| Capital share (Cobb–Douglas) | $\alpha$   | 0.35             | –                         |
| TFP multiplier               | $\bar{Z}$  | $10^{1-\alpha}$  | Normalization             |
| Persistence of log TFP       | $\rho_Z$   | 0.90             | –                         |
| Std. dev. of TFP shock       | $\sigma_Z$ | 0.0152           | $\text{std}(Z_t) = 3.5\%$ |
| Depreciation rate            | $\delta$   | 0.10             | –                         |

### 1.4.1 Data Source

For data, I use the Survey of Consumer Finances (SCF) 2019 release (Board of Governors of the Federal Reserve System 2019). The time reflected in this sample is roughly co-incident with the announcement of the Biden Administration’s planned policy, so this should represent the most accurate data to study.

The SCF oversamples high-wealth and high-income households, so I truncate the top 10% of the income distribution. Since samples are relatively small when segmented by age and loan-to-income rankings, the income and student debt life-cycle profiles are not smooth. To aid with computational efficacy, I choose to smooth these profiles before feeding them in to my model. For labor efficiency units, I choose to smooth with a cubic polynomial fit to reflect the well-established hump shape of the income profile. For student debt balances, I use a quadratic polynomial fit and for student debt payments I use a linear polynomial fit.




Figure 1.1: Estimation of life-cycle profile of income, student debt balances, and student debt payments. Also plotted: cubic, quadratic, and linear polynomials respectively to smooth the profiles for computation. Source: SCF 2019.

Since student debt balances are sometimes borne by the *parents* of the college student requesting financing, student debt balances in the data stay high through the life cycle. To reflect this realistic family structure while staying agnostic to family linkages, I normalize all life-cycle profiles from the SCF by household size, which is reported in the SCF.

#### 1.4.2 Estimating Student Debt Balances After Forgiveness

A central challenge in modeling the Biden Administration's forgiveness proposal is that the policy has not been implemented, making post-forgiveness debt balances and payment schedules unobservable in the data. This necessitates assumptions about how forgiveness would affect the life-cycle profiles of debt balances and required payments for each household type. This section describes my approach to constructing these counterfactual profiles.

The Biden Administration's plan calls for forgiveness of up to \$10,000 for earners below \$125,000 (single) or \$250,000 (married), with up to \$20,000 in forgiveness available to Pell Grant recipients. Since I aggregate borrowers into three broad types (non-borrowers, low loan-to-income borrowers, and high loan-to-income borrowers), I cannot simply reduce all balances by a uniform \$10,000. Instead, the amount of forgiveness received by each type at each age represents the *expected* forgiveness conditional on age, type, income, and Pell Grant recipient status. To construct this expected forgiveness, I use the SCF 2019 microdata to compute the average forgiveness that would accrue to households of each age and loan-to-income type, accounting for both the income eligibility thresholds and the differential treatment of Pell Grant recipients.

Figure 1.2 (top panel) shows the resulting life-cycle profile of average forgiveness amounts by type. Several features are worth noting. First, forgiveness amounts decline with age for borrowers, reflecting two forces: (a) older borrowers tend to have lower remaining balances due to years of repayment, and (b) older borrowers are more likely to have incomes exceeding the \$125,000 threshold, rendering them ineligible for relief. Second, high loan-to-income borrowers (type  $j = 2$ ) receive higher average forgiveness than low loan-to-income borrowers (type  $j = 1$ ) at young ages, since their balances are higher and a

larger share are Pell Grant recipients who qualify for the \$20,000 tier. Third, the forgiveness amounts are modest relative to outstanding balances, particularly for older borrowers, implying that forgiveness provides partial rather than complete relief.

Given the post-forgiveness debt balances, I must next determine the implied payment schedules. Since the counterfactual payment schedules are unobserved, I adopt a parsimonious approach that maintains a constant payment-to-balance ratio across the before- and after-forgiveness policy environments. Formally, let  $d_{i,j}$  denote the baseline (no-forgiveness) debt balance for age  $i$  and type  $j$ , let  $a_{i,j}$  denote the baseline payment, and let  $d_{i,j}^{\text{forg}}$  denote the post-forgiveness balance. I assume the post-forgiveness payment  $a_{i,j}^{\text{forg}}$  satisfies:

$$\frac{a_{i,j}^{\text{forg}}}{d_{i,j}^{\text{forg}}} = \frac{a_{i,j}}{d_{i,j}}. \quad (1.19)$$

This assumption is both model-free and empirically plausible. If household payments follow a standard amortization formula (as is typical for federal student loans under standard repayment plans), reducing the principal balance by a constant proportion mechanically reduces required payments by the same proportion, holding the interest rate and remaining maturity fixed. This approach avoids imposing additional structure on repayment behavior and abstracts from heterogeneity in repayment plans (e.g., income-driven repayment vs. standard repayment), which would require strong assumptions about which plans borrowers select and how those selections might change post-forgiveness.

Figure 1.2 (middle and bottom panels) displays the baseline and post-forgiveness life-cycle profiles of debt balances and payments for each borrower type. The middle panel shows that forgiveness generates a discrete downward shift in balances at all ages, with the magnitude of the shift declining with age as described above. The bottom panel shows that payments decline proportionally, preserving the life-cycle shape of the payment profile while reducing its level. Importantly, even after forgiveness, high loan-to-income borrowers (type  $j = 2$ ) face substantially higher debt burdens and payments than low loan-to-income borrowers (type  $j = 1$ ), indicating that forgiveness does not eliminate heterogeneity in financial distress across borrower types.

I acknowledge that this constant payment-to-balance ratio assumption is a simplification. In reality, some borrowers may choose to maintain their pre-forgiveness payment levels to pay down principal faster, while others may be enrolled in income-driven repayment plans where payments depend on income rather than balance. However, absent detailed micro-data on repayment plan enrollment and borrower optimization over repayment strategies—which are beyond the scope of this paper—the constant ratio assumption provides a transparent, model-free benchmark. Future work could refine this approach by explicitly modeling endogenous repayment plan choice or by incorporating richer micro-data on actual post-forgiveness payment behavior (should such data become available). For the purposes of this paper, the constant ratio assumption allows me to focus on the central question of whether forgiveness enables wealth accumulation, rather than on the mechanics of debt amortization.




Figure 1.2: Estimated profiles of student debt balance and payments after forgiveness according to the White House plan: \$10,000 in forgiveness for income below \$125,000. Older borrowers receive (on average) lower forgiveness because of (a) lower debt balances or (b) higher incomes. Source: SCF 2019.

## 1.5 Solution Method: FINNs

The general structure of the algorithm used to solve the fully stochastic model is inspired by Azinovic, Gaegauf, and Scheidegger [2019](#) and Azinovic and Zemlicka [2024](#), though with a few differences detailed below. For the most up-to-date primer on deep learning for economics, including first-principles building of (deep) neural networks, see Fernández-Villaverde [2025](#).

Following the physics-informed neural networks (PINNs) approach of Raissi, Perdikaris, and Karniadakis ([2019](#)), which embeds partial differential equations directly into the neural network loss function, I develop what I term finance-informed neural networks (FINNs). Rather than the PDEs governing a *physical* system, FINNs incorporate the equations governing the *financial* model—Euler equations, market clearing conditions, and other terms discussed below—as penalty terms in the loss function. This approach ensures that the trained FINN respects the economic structure of the problem while approximating the equilibrium policy functions.

### 1.5.1 Applications of Domain Knowledge

Naïve machine learning methods excel at finding solutions that adhere to the minimization problem of the loss function as specified by the user. The upside of this means that economists need not spend significant time writing the low-level algorithm needed for training the FINN. The downside is that the landscape may admit low-error pathological ‘solutions’ that evaluate to a very small error but are economically implausible. Examples of this that I encountered in training often include:

- constant or near-constant predictions that fail to vary with the underlying shock process in any way;
- predictions that vary only in measure isomorphism to the underlying shock process;
- prices approaching zero and quantities approaching infinity;

- and other pathological outcomes.

Especially in the stochastic overlapping generations environment where Euler equation iteration is not necessarily a contraction mapping, these outcomes must be monitored.

To avoid the pathological predictions of the type discussed in the preceding paragraph, I employ the following custom strategies to augment the off-the-shelf routines available in coding software.

1. *Using the loss function to discipline economically plausible outcomes.* The formal statement of the loss function is below, but it can be understood as comprising of three parts:

- (a) The mean-square-error (across households and time) of the Fisher-Burmeister transform between the Euler equation residual and forward capital;
- (b) The mean-square-error (across time) of the aggregate feasibility constraint. This ensures the pathological solution of  $c_{i,j,t} \rightarrow 0$  for any  $i, j, t$  is not found; and
- (c) A penalty term inversely related to the variance (across time) of the aggregate capital stock. This ensures the pathological solution of households refusing to adjust their capital accumulation across the business cycle is not viable.

2. *Applying the Fisher-Burmeister transformation to eliminate the computation of Lagrange multipliers.* Since the agents in the model are constrained to not sell capital short, the Euler equation may not bind endogenously for all agents. As formulated in the formal model earlier in this paper, this entails solving for Lagrange multipliers of each agent, effectively doubling the required state space. To avoid this, Azinovic and Zemlicka (2024) applies the Fisher-Burmeister transform as follows:

$$\Psi^{\text{FB}}(a, b) = a + b - \sqrt{a^2 + b^2} \quad (1.20)$$

which embeds the KKT conditions as  $a, b \geq 0$  and  $ab = 0 \iff \Psi(a, b) = 0$ . However, since the FINN takes derivatives of equilibrium conditions, the Fisher-Burmeister equation becomes

unstable as  $a, b \rightarrow 0$  jointly: the derivative of the square-root approaches positive infinity. To stabilize this, I modify the Fisher-Burmeister equation to:

$$\Psi(a, b; \lambda_{\text{FB}}, \epsilon_{\text{FB}}) = \lambda_{\text{FB}} \left( a + b - \frac{a^2 + b^2}{\sqrt{a^2 + b^2 + \epsilon_{\text{FB}}^2}} \right) + (1 - \lambda_{\text{FB}})a^+b^+ \quad (\text{I.21})$$

where for this model I set  $\lambda_{\text{FB}} = 0.8$  and  $\epsilon_{\text{FB}} = 10^{-3}$ .

3. *Time  $t = 0$  initialization near or on the equilibrium manifold.* Since—as stated previously—iteration on the Euler equations in stochastic overlapping generations models is *not* a contraction mapping, the starting guess is very important to finding equilibria. Some approaches entail a pre-training phases of an easier model, such as with fewer assets or no aggregate risk. Since the model in this paper is already quite simple, I do not need to employ these methods. Instead, after each simulation step of drawing the ergodic time series, I will save and store  $\bar{\mathcal{X}}$  as the mean across time of the FINN inputs. Since the FINN is learning the global policy functions, these values can only be expected to evaluate correctly on admissible equilibrium values. In the subsequent time series iteration, I will initialize endogenous state variables at time  $t = 0$  based on the values in  $\bar{\mathcal{X}}$ .
4. *Alleviating the sequential bottleneck in data generation for training.* Unlike in many machine learning environments, the data on which the algorithm trains for this model is itself generated by the model. In particular, the policy function predicts forward capital in this model. In order to generate training data, I must simulate outcomes of the aggregate risk process along with predictions for forward capital, which become the capital in the next period. Usual sample sizes for models like this are for around  $T = 10000$  periods. Rather than draw one block of 10000 periods, I instead draw 100 blocks of 100 periods *in parallel*, enormously parallelizing the simulation pass for each loop in the training algorithm.

### 1.5.2 FINN Architecture and Training

Theory indicates that the state variable required for computation of the model in this paper is  $\mathcal{X}_t = (\{k_{i,j,t}\}_{i,j}, Z_t)$ , and the predictions of the FINN will be  $\mathcal{Y}_t = (\{k_{i+1,j,t+1}\}_{i,j})$ . The FINN architecture consists of two hidden layers, each with 100 neurons. With  $I = 60$  periods of life and  $J = 3$  types, the input layer has dimension  $3 \times 59 + 1 = 178$  (capital holdings for all agents plus the aggregate productivity shock  $Z_t$ ), and the output layer has dimension  $3 \times 59 = 177$  (next-period capital holdings for all agents). I apply the hyperbolic tangent activation function  $\tanh(\cdot)$  to the hidden layers, which provides smooth, differentiable nonlinearity necessary for backpropagation through the economic equilibrium conditions. Since capital holdings must be nonnegative, I apply the softplus activation function  $\text{softplus}(x) = \log(1 + e^x)$  to the output layer, which smoothly enforces the nonnegativity constraint while remaining differentiable everywhere.

To ensure stable training and keep inputs in the active region of the  $\tanh(\cdot)$  activation function (where gradients are meaningfully nonzero), I normalize all inputs by dividing by  $1 + \overline{\mathcal{X}}$ , where  $\overline{\mathcal{X}}$  is the running mean of the state variable across training iterations. This normalization prevents vanishing gradients and avoids division by zero. The network is trained using the Adam optimizer (Kingma and Ba 2014) with a small learning rate of  $10^{-6}$  to ensure stability in the presence of the complex, nonlinear economic constraints embedded in the loss function. Training proceeds for 10,000 episodes, with minibatches of size 200. Crucially, the simulated training data is redrawn after each epoch, ensuring that the network learns from fresh realizations of the stochastic equilibrium rather than overfitting to a single simulation path. Across all solved models, the loss function evaluates to approximately  $2 \times 10^{-6}$  and the Fisher-Burmeister transform of the Euler equation residual averages approximately 0.15% per household, indicating that the trained networks accurately satisfy the equilibrium conditions.

Expectations in the Euler equations and forecasts of future state variables are computed using Gauss-Hermite quadrature with 15 nodes. Table 1.2 reports the computational environment and runtime statistics. Training each model to convergence requires approximately 45 minutes on consumer-grade hard-

ware. Unlike machine learning applications that require large pre-existing datasets, the FINN approach generates its own training data through simulation. This means that the memory bottleneck is the FINN size rather than dataset size, allowing the entire training process to fit comfortably within the VRAM limits of consumer GPUs. This demonstrates the computational accessibility of the FINN approach for large-scale stochastic OLG models without requiring specialized high-performance computing infrastructure. I simulate the training data on the CPU (which is faster at this sort of task) and train the FINN (constructing economic quantities and computing loss function, evaluating the parameters step) on the GPU.

Table 1.2: Computational Environment and Performance

| <b>Hardware</b>                     |                                           |
|-------------------------------------|-------------------------------------------|
| Processor                           | 12th Gen Intel i9-12900KF (24) @ 5.100GHz |
| GPU                                 | NVIDIA GeForce RTX 3080 with 8GB of VRAM  |
| RAM                                 | 32 GB                                     |
| <b>Performance</b>                  |                                           |
| Approximate training time per model | 45 minutes                                |
| Training episodes                   | 10,000                                    |
| Minibatch size                      | 200                                       |
| Quadrature nodes                    | 15 (Gauss-Hermite)                        |
| Approximate Euler residual          | 0.15%                                     |
| Approximate Loss function value     | $2 \times 10^{-6}$                        |

We can formally define the FINN parameterized by  $\Theta$  as a function that maps the state variable to the policy functions:

$$V_{\Theta}(\cdot) : \mathcal{X}_t \mapsto \mathcal{Y}_t$$

where  $\mathcal{X}_t = (\{k_{i,j,t}\}_{i,j}, Z_t)$  represents the complete state of the economy (all agents' capital holdings and the aggregate productivity shock) and  $\mathcal{Y}_t = (\{k_{i+1,j,t+1}\}_{i,j})$  represents the equilibrium policy functions (all agents' optimal forward capital choices). The neural network  $V_{\Theta}$  thus approximates the time-homogeneous Markov equilibrium policy functions that characterize the recursive competitive equilibrium defined earlier.

Critically, given the state variable  $\mathcal{X}_t$  and the FINN's predicted policy functions  $\mathcal{Y}_t$ , I can recover all economic quantities needed to evaluate equilibrium conditions using the budget constraints, market clearing conditions, and firm optimality conditions from the formal model. Specifically: aggregate capital  $K_t = \sum_{i,j} k_{i,j,t}$  and aggregate labor  $L_t = \sum_{i,j} \ell_{i,j}$  determine factor prices via the firm's optimality conditions  $(w_t, r_t)$ ; these prices combined with the exogenous efficiency units  $\ell_{i,j}$  and debt payment schedules  $a_{i,j}$  determine household incomes and the government budget constraint determines the tax rate  $\tau_t$ ; and finally the household budget constraints then back out consumption  $c_{i,j,t}$  and investment  $i_{i,j,t}$  for all agents. This closed-form recovery of all equilibrium objects from  $(\mathcal{X}_t, \mathcal{Y}_t)$  is what allows the loss function to be evaluated purely as a function of the neural network's inputs and outputs, without requiring additional solution steps within each training iteration.

The loss function can be evaluated as

$$\begin{aligned} \mathcal{L}_\Theta(\{\mathcal{X}_t, \mathcal{Y}_t\}_t) = \log \left( 1 + \right. \\ \left. \frac{1}{TJ(I-1)} \sum_{i,j,t} \Psi \left( \frac{(u')^{-1} (\beta \mathbb{E}[u'(c_{i+1,j,t+1})(1+r_{t+1}-\delta)])}{c_{i,j,t}} - 1, k_{i+1,j,t+1} \right)^2 \right. \\ \left. + \frac{100}{T} \sum_t \left( \frac{\sum_{i,j} (c_{i,j,t} + i_{i,j,t} + d_{0,j})}{Y_t} - 1 \right)^2 + \frac{10^{-3}}{\text{Var}[\sum_{i,j} k_{i+1,j,t+1}]} \right) \quad (1.22) \end{aligned}$$

where the three terms inside the logarithm correspond to: (1) the mean-square-error of the Fisher-Burmeister-transformed Euler equation residuals across all households and time periods, ensuring that the first-order conditions for household optimization are satisfied; (2) the mean-square-error of the aggregate resource constraint, ensuring that consumption, investment, and government spending sum to output and thereby ruling out pathological solutions where consumption approaches zero; and (3) an inverse variance penalty on aggregate capital, ensuring that the policy functions respond meaningfully to aggregate shocks rather than producing near-constant predictions. Together, these three components discipline the neural network to find economically plausible solutions that satisfy household and firm optimizations, market clearing, and balance the government budget constraint.

The outer logarithmic transformation  $\log(1 + \cdot)$  serves an important numerical stability purpose during training. In early training iterations when the network parameters are far from equilibrium, the raw loss terms inside the parentheses can be very large, potentially causing gradient explosion and numerical overflow. The logarithmic transformation compresses these large values, preventing the loss from growing unboundedly and ensuring stable gradient-based optimization. Conversely, when the network has converged and the raw loss is very small (near zero), the logarithm satisfies  $\log(1 + x) \approx x$  for small  $x$ , meaning that the transformation becomes approximately linear and does not distort the loss landscape near the optimum. This design ensures stable training throughout the optimization process while preserving sensitivity to small equilibrium violations once the network approaches convergence.

Finally, we have all the pieces in place to define the formal problem solved by the FINN training algorithm:

$$\Theta^* = \arg \min_{\Theta} \mathcal{L}_{\Theta}(\{\mathcal{X}_t, \mathcal{Y}_t\}_t) \quad (1.23)$$

where the set of neural network parameters  $\Theta^*$  will be discovered through the training routine specified in code, usually a modification of stochastic gradient descent such as the Adam optimizer. The trained FINN  $V_{\Theta^*}$  then provides the solutions to the economic model: for any state  $\mathcal{X}_t$ , the equilibrium policy functions are given by  $\mathcal{Y}_t = V_{\Theta^*}(\mathcal{X}_t)$ , and all other equilibrium objects (consumption, prices, taxes) can be recovered algebraically from the budget constraints and market clearing conditions as described above. By minimizing the loss function—which embeds the economic equilibrium conditions—the FINN training algorithm effectively solves for the recursive Markov equilibrium of the stochastic overlapping generations model.

## 1.6 Results

To analyze the results of the Biden Administration student debt forgiveness policy proposal, I first solve the baseline economy. I then solve the economy under the proposed \$10,000 forgiveness policy. Additionally, I solve the model according to two counterfactual policies to better understand the underlying economic mechanisms. In the first, I solve the model under full student loan forgiveness—where student debt is never repaid. In the final model, I solve the model under full student loan forgiveness, but holding the tax obligation of non-borrowers pinned at the level of the baseline economy. This experiment should isolate pure general equilibrium effects of the forgiveness policy on the non-borrowers. For each of the four policy environments, I simulate the economy for 100,000 periods to ensure precise estimates of welfare and other equilibrium statistics.

The tax experiment reveals that almost all welfare loss borne by the non-borrowers under the partial \$10,000 forgiveness and full forgiveness is a result of a higher tax obligation (Figure 1.3).

Table 1.3: Consumption-Equivalent Welfare by Group Across Economies

|          | Baseline | \$10,000 | Full Forgiveness | Tax Experiment |
|----------|----------|----------|------------------|----------------|
| No Debt  | –        | -1.15%   | -1.77%           | -0.10%         |
| Low LTI  | –        | +1.46%   | +1.85%           | +0.47%         |
| High LTI | –        | +0.71%   | +5.08%           | +3.70%         |




Figure 1.3: Consumption-Equivalent Welfare

For borrowers with student debt, particularly the high loan-to-income types, student loan forgiveness does *not* lead to earlier wealth accumulation (Figure 1.4). In other words, the life-cycle motives for consumption smoothing overpower the deleveraging effect of student loan forgiveness.




Figure 1.4: Expected Life-Cycle Capital Profiles Across Economies

In contrast, the forgiveness of student loans *does* increase early-age consumption by financially constrained households (Figure 1.5). This is despite the quite aggressive parameterization of the household utility function, with  $\gamma = 3$ . At this level of risk aversion, precautionary savings motives should be quite intense so barring financial constraints, households should be eager to build savings early in life. Their unwillingness to meaningfully alter their savings behavior is evidence that the student loan obligations were not the cause of underinvestment in retirement savings (or entrepreneurial enterprises) for young borrowers of student loans.




Figure 1.5: Expected Life-Cycle Consumption Profiles Across Economies

Since household savings (in the form of the accumulation of productive capital) is not meaningfully influenced by the student loan forgiveness environment, it is not surprising that real outcomes do not vary much: production, wages, aggregate capital, and returns barely vary across these policy environments (Tables 1.4 and 1.5).

This lack of general equilibrium price effects stands in contrast to the theoretical predictions from the literature on overlapping generations asset pricing. As argued by Constantinides, Donaldson, and Mehra (2002), shifting financial wealth toward younger, more financially constrained households should alter the characteristics of the marginal investor, thereby affecting equilibrium asset prices. In particular, one might expect student loan forgiveness to decrease returns as young households, freed from debt obligations, enter asset markets more aggressively. However, this mechanism fails to materialize in quantitatively meaningful ways precisely because the forgiveness, dwarfed by life cycle effects, does not translate

into higher capital accumulation by borrowers.

The muted response of aggregate capital stock to the forgiveness policy also implies limited feedback effects on wages and returns. Since borrowers choose to consume rather than invest the transfer, the aggregate capital stock remains nearly constant across policy environments. Consequently, the marginal product of labor (and thus wages) and the marginal product of capital (and thus returns) are virtually unchanged. This finding undermines one potential justification for student loan forgiveness from a non-borrower perspective: that the policy might generate positive spillovers through higher aggregate productivity or wages due to deleveraged, more productive entrepreneurs and workers. Instead, the results suggest the primary effect is a pure fiscal transfer with minimal real economic ramifications.

Table 1.4: Asset Pricing Outcomes Across Economies

|                    | Baseline | \$10,000 | Full Forgiveness | Tax Experiment |
|--------------------|----------|----------|------------------|----------------|
| Average return (%) | 12.86    | 12.98    | 12.89            | 12.95          |
|                    | –        | +12 bps  | +3 bps           | +8 bps         |
| Volatility (%)     | 0.82     | 0.78     | 0.77             | 0.77           |
|                    | –        | -4 bps   | -5 bps           | -5 bps         |

Table 1.5: Real Outcomes Across Economies

|                       | Baseline | \$10,000 | Full Forgiveness | Tax Experiment |
|-----------------------|----------|----------|------------------|----------------|
| Output                | 12.599   | 12.559   | 12.605           | 12.587         |
|                       | –        | -0.3%    | +0.0%            | -0.1%          |
| Output Volatility (%) | 5.620    | 5.468    | 5.411            | 5.489          |
|                       | –        | -15 bps  | -21 bps          | -13 bps        |
| Capital               | 19.322   | 19.157   | 19.301           | 19.227         |
|                       | –        | -0.9%    | -0.1%            | -0.5%          |
| Wage                  | 8.190    | 8.163    | 8.193            | 8.181          |
|                       | –        | -0.3%    | +0.0%            | -0.1%          |

One avenue left to explore is the avenue of generational risk sharing. The federal government's dominant role in student loan issuance—accounting for approximately 89% of outstanding balances—can be understood as the government acting as lender of last resort to young borrowers who may lack credit histories and collateral. Private credit markets would either decline to lend to such borrowers or would charge prohibitively high interest rates to compensate for default risk and adverse selection. By offer-

ing federally-backed loans at statutorily-set interest rates, the government plausibly engages in a form of intergenerational risk sharing that private markets cannot provide: spreading the risk of uncertain post-graduation earnings across taxpayers rather than concentrating it on individual young borrowers. This raises a natural question: does student loan forgiveness enhance this risk-sharing function by further stabilizing consumption for financially constrained households facing aggregate shocks? In stochastic overlapping generations economies, households are unable to perfectly insure themselves against risk by trading contingent claims with other households due to the restricted market participation: some households are no longer alive and not yet born (Cass and Shell 1983). If forgiveness reduces consumption volatility—particularly for highly-leveraged borrowers most exposed to income risk—it may improve welfare through this risk-sharing channel even in the absence of effects on capital accumulation or aggregate output.




Figure 1.6: Life-Cycle Consumption Risk Profiles Across Economies

My results show that the answer is yes, though the magnitude is small. Consumption volatility decreases not only for borrowers but also for non-borrowers, despite the latter facing higher tax obligations (Figure 1.6). Student loan forgiveness acts as a form of government-provided intergenerational risk sharing. The mechanism operates through general equilibrium stabilization. Prior to forgiveness, borrowers faced fixed debt obligations that forced sharp consumption cuts during economic downturns, amplifying aggregate demand volatility. By deleveraging these households, forgiveness stabilizes their consumption responses to aggregate shocks, which in turn stabilizes equilibrium wages and returns. Non-borrowers benefit from this stabilization even though they bear the fiscal cost through higher taxes. In effect, the government partially completes markets by redistributing aggregate risk away from the financially constrained young and spreading it across all generations through the tax system.

To isolate the welfare effects attributable purely to changes in consumption risk, I compute a risk-adjusted consumption-equivalent welfare (CEW) measure. This is constructed by first computing the standard CEW across the full stochastic equilibrium, then computing the CEW based only on the mean consumption levels (abstracting from risk), and taking the difference between these two measures. The resulting risk component captures the welfare change due solely to altered exposure to aggregate shocks, separately from the level effects of the transfer.

The welfare gains from the intergenerational risk sharing mechanism are very small and do not offset the welfare losses on non-borrowers from fiscal effects (Figure 1.7).




Figure 1.7: Consumption-Equivalent Welfare: Risk Component

## 1.7 Conclusion

This paper evaluates the Biden Administration’s proposed student loan forgiveness policy in a general equilibrium framework to understand its redistributive effects across borrowers and non-borrowers. Motivated by the policy’s stated goal of enabling wealth accumulation through reduced debt burdens, I study whether student loan forgiveness generates meaningful changes in household savings behavior and asset prices, or whether it primarily represents a fiscal transfer. The policy environment examined involves up to \$10,000 in debt cancellation for eligible borrowers, with an additional \$10,000 available to Pell Grant recipients, subject to income thresholds.

To capture the intergenerational and asset pricing effects of student loan forgiveness, I develop a fully stochastic overlapping generations model with 60 periods of life and three household types differentiated by their student debt-to-income ratios. The model features aggregate productivity risk, endogenous capital accumulation, and government-issued student loans financed through a balanced budget. By incorporating long-lived agents in a stochastic environment, the model generates uninsurable generational risk and allows for rich interactions between household debt obligations, life-cycle savings decisions, and equilibrium asset prices. This framework is essential for studying policies that transfer resources across age cohorts and may alter the characteristics of the marginal investor in asset markets.

The model is calibrated to the 2019 Survey of Consumer Finances, matching life-cycle profiles of income, student debt balances, and repayment schedules across household types. Standard macroeconomic parameters governing preferences, technology, and aggregate risk are set in line with the empirical literature. The calibration ensures that the baseline economy replicates key features of the U.S. economy, including the distribution of student debt across age and income and realistic capital accumulation patterns over the life cycle.

I solve the model using finance-informed neural networks (FINNs), a computational method that embeds economic equilibrium conditions directly into the loss function of a deep neural network. This approach allows me to compute global policy functions for the high-dimensional state space without

relying on dimensionality reduction or local approximation techniques. Each of the four policy environments is simulated for 100,000 periods, and the trained neural networks achieve Euler equation residuals averaging 0.15% across all households, indicating accurate solutions to the equilibrium conditions.

The central finding is that in contrast with its stated goal, the Biden Administration student loan forgiveness program generates minimal real economic effects. Despite reducing debt burdens, borrowers increase consumption rather than capital accumulation, leaving aggregate savings and production virtually unchanged across policy environments. Consequently, equilibrium wages, returns, and asset prices exhibit negligible variation. For non-borrowers, the policy delivers welfare losses driven almost entirely by higher tax obligations needed to finance the forgiveness, with no offsetting gains from general equilibrium spillovers. The small positive welfare effect from reduced consumption risk across all household types does not offset these fiscal costs. These results contradict the policy's stated objective of promoting wealth accumulation and suggest that student loan obligations were not the binding constraint on retirement savings for young borrowers. The findings underscore that the Biden Administration's forgiveness proposal functions primarily as a fiscal transfer rather than a mechanism to unlock productive investment or alter macroeconomic outcomes.

It is important to emphasize that these findings do not address the normative question of *whether* student loan forgiveness should be pursued. There may be compelling equity-based arguments for providing relief to the most financially distressed borrowers—particularly those with high debt-to-income ratios and limited earnings prospects—on grounds of fairness, compassion, or as a form of social insurance against adverse education-labor market outcomes. My analysis demonstrates only that the stated *mechanism*—unlocking productive investment and wealth accumulation—does not operate as advertised in general equilibrium. Policymakers concerned with borrower welfare may find these results informative for designing more effective interventions, but the case for or against debt relief ultimately rests on political judgments beyond the scope of positive economic analysis.

More fundamentally, the one-time nature of the Biden Administration's forgiveness proposal fails to address the structural forces that generated the student debt crisis in the first place. The policy provides

temporary relief to current borrowers while leaving intact the dual mechanisms accelerating student debt accumulation: unbounded government loan issuance and persistently rising tuition prices charged by colleges and universities. Without reforms to these underlying mechanisms, future cohorts of students will accumulate debt burdens similar to or exceeding those of current borrowers, necessitating repeated forgiveness interventions and perpetuating moral hazard on both the demand and supply sides of the higher education market. A more durable solution would require institutional changes that align the incentives of government lenders and educational institutions with the long-run financial well-being of students. Future research could fruitfully explore bargaining models or contract design frameworks in which the government and universities negotiate over tuition pricing, loan terms, and risk-sharing arrangements—for instance, conditioning institutional funding on post-graduation earnings outcomes or requiring universities to retain skin-in-the-game through partial loan guarantees. Such mechanisms could mitigate the moral hazard inherent in the current system while preserving broad access to higher education financing.

## Chapter 2: Real and Asset Pricing Effects of Employer Retirement Matching

### 2.1 Introduction

A key feature of modern retirement planning (often in the form of defined-contribution plans like 401(k) or 403(b)) is employer matching contributions, whereby firms commit to match employee retirement contributions up to a specified threshold—typically some fraction of salary. Contributions to these plans comprise over \$500 billion annually in the United States alone, with approximately two-thirds of private-sector workers having access to matching programs. The generosity of matching schedules has grown substantially over the past two decades, reflecting its increasing centrality to retirement security and household wealth accumulation. This practice augments workers’ retirement savings in the form of equity holdings. By subsidizing equity purchases, households receiving the match may tolerate lower returns since the total return remains attractive inclusive of the match. Since households own the firm through their equity holdings, this reduction in returns represents a decline in the firm’s cost of equity capital—the rate at which the firm discounts future dividend streams. Facing a lower cost of capital, the firm optimally increases investment in physical capital, potentially raising aggregate output, wages, and welfare. The central question of this paper is whether this general equilibrium pricing mechanism generates meaningful macroeconomic effects on firm investment, asset prices, and economic output, or whether employer matching has negligible effects on real outcomes beyond the direct subsidy to participating workers.

This paper develops a stochastic overlapping generations model paired with a neoclassical model of firm investment to study these questions. A central methodological contribution is the integration of Samuelson (1958)-style overlapping generations on the household side with the q-theory of investment (Tobin 1969; Jorgenson 1963) on the firm side, allowing retirement policy to affect corporate capital accumulation through endogenous equity pricing. Households live for multiple periods, choosing consumption and equity holdings to finance retirement, and receive employer matching contributions according to a realistic matching schedule (a fraction  $\psi$  of contributions up to a cap  $\phi$  times income). The representative firm faces convex adjustment costs in the tradition of Hayashi (1982) and makes dynamic investment decisions, discounting future dividends at the endogenous stochastic discount factor (SDF) implied by equity holders. The rich link between households' private savings behavior and firm real outcomes via the firm using the endogenous SDF from household Euler equations provides a rich environment to study the link between corporate finance and household retirement financing. The model features aggregate productivity risk, allowing me to study how matching affects not only the level but also the risk properties of returns and investment.

I find along the lines of the mechanism previewed above, employer matching generates substantial general equilibrium effects beyond simple redistribution. Matching increases aggregate savings by subsidizing workers' equity accumulation, which in turn reduces equilibrium equity returns and increases the SDF. The resulting lower cost of capital stimulates firm investment, raising the aggregate capital stock, output, and wages. These results stand in contrast to a partial equilibrium view in which matching merely transfers resources from shareholders to workers. Instead, the policy reshapes the investment opportunity set available to firms, aligning this paper with recent work showing that innovations in retirement finance can affect real corporate behavior through general equilibrium price effects (A. Zhang 2022).

My model assumes that employer matching contributions are financed entirely out of dividends rather than being passed through to workers via lower wages. In reality, the incidence of matching costs may be partially borne by workers, or determined through some form of Nash bargaining between firms and employees. However, this simplification does not undermine the core mechanism. Regardless of whether

matching is financed by reduced dividends, reduced wages, or some combination thereof, the subsidy to workers' equity purchases remains operative. Employer matching acts as a multiplier on equity returns from the household perspective: by subsidizing equity purchases, matching effectively increases the return households receive on their own contributions, inducing them to save more and hold more equity. This increased equity demand raises the SDF—households become more patient in equilibrium—which reduces equilibrium equity returns and lowers the firm's cost of capital. Crucially, this channel operates through the pricing kernel of the marginal investor, not through mechanical effects of dividend payments. The key insight—that employer matching affects corporate capital accumulation through endogenous equity pricing and the SDF—operates independently of the specific incidence of matching costs and represents a first-order general equilibrium mechanism even in the absence of tax distortions or wage bargaining.

### 2.1.1 Institutional Setting of Firm Retirement Contributions

The shift from defined-benefit (DB) to defined-contribution (DC) retirement plans accelerated following the Employee Retirement Income Security Act of 1974, which created the regulatory framework for employer-sponsored retirement plans (Kruse 1995). Subsequent legislation—including the Economic Growth and Tax Relief Reconciliation Act of 2001, which raised contribution limits, and more recently the SECURE Acts I and II (Reuter 2024)—has further expanded tax incentives and flexibility for DC plans. Arnoud et al. (2021) document that employer matching schedules became increasingly generous between 2003 and 2017, with the combined employer-employee contribution rate rising by approximately one percentage point. As of recent data, roughly two-thirds of private-sector workers have access to employer-sponsored DC plans, with the vast majority of these plans offering some form of employer matching. Given this widespread adoption and the increasing generosity of matching contributions, understanding the general equilibrium effects of employer matching on asset prices, capital accumulation, and macroeconomic outcomes has become imperative. While an economist could spend their entire career unpacking the rich institutional details, tax incentives, and policy nuances in this space, the central pricing mechanism I study—how matching subsidies affect household equity demand and firm invest-

ment through the endogenous SDF—is robust across these institutional variations.

For firms, employer matching contributions offer several advantages. First, matching contributions are tax-deductible as compensation expense, reducing the firm’s taxable income. Second, employer contributions are exempt from payroll taxes, providing additional tax savings. Third, DC plans impose lower administrative costs than DB plans by shifting investment risk and longevity risk from the firm to employees themselves. Fourth, firms retain flexibility to suspend or reduce matching contributions during economic downturns, unlike the fixed obligations associated with DB pension promises.

For employees, the tax advantages are equally compelling. Employee contributions to 401(k) plans are made on a pre-tax basis, reducing current taxable income. Investment returns within the account grow tax-deferred until withdrawal at retirement. Employer matching contributions represent additional compensation that is not taxed until withdrawal.

Despite these tax advantages, the model presented in this paper abstracts from explicit taxation—of firm income, payrolls, capital, and labor—to focus on the general equilibrium mechanisms linking retirement policy to firm investment. This modeling choice prioritizes parsimony and isolates the core economic channel: employer matching subsidizes equity demand, which affects equity prices and firm investment through the cost of capital. Introducing taxation would likely amplify the effects found in this paper, as the tax deductibility of employer contributions would further reduce the effective cost of matching for firms, potentially increasing the equilibrium level of matching and magnifying the impact on capital accumulation.

### 2.1.2 Preview of Results

To understand the mechanisms of study in this paper, we can start with the central asset pricing relationship:

$$1 = \mathbb{E} \left[ \underbrace{\frac{\beta u'(c')}{u'(c)}}_{\text{SDF}} \cdot \left( 1 + \frac{\partial m}{\partial s} \right) \cdot \underbrace{\frac{p' + d'}{p}}_{R'} \right] \quad (2.1)$$

The key insight is that employer matching directly enters the SDF through the term  $(1 + \frac{\partial m}{\partial s})$ , which captures the marginal subsidy to equity purchases. When matching intensity increases, households face a stronger incentive to save in order to capture the employer match—leaving matching dollars on the table is equivalent to forfeiting free money. The desire not to leave money on the table leads households to purchase more equity earlier in the life cycle, tolerating lower market returns because their effective returns inclusive of the match remain attractive. Since the asset pricing equation must hold in equilibrium, an increase in the SDF forces equilibrium returns  $R'$  to fall to maintain the equality. This mechanism operates independently of any mechanical dividend effects: the return decline reflects the endogenous adjustment of the pricing kernel, not a reduction in cash flows. The firm, observing this lower cost of capital, optimally responds by increasing investment in physical capital.

To establish the core mechanism transparently, I first derive closed-form analytical results in a simplified two-period model without aggregate risk. In this parsimonious setting, I show that employer matching unambiguously increases the SDF, reduces equilibrium equity returns, and stimulates firm capital investment. The analytical results reveal the key economic forces at work: matching subsidizes equity purchases, inducing households to tolerate lower returns, raising the SDF. Facing a higher SDF, the firm discounts future dividends less heavily and optimally increases investment in physical capital and therefore production. These closed-form results provide clean intuition for the mechanism and demonstrate that the qualitative effects are robust features of the model structure rather than artifacts of specific functional forms or calibration choices.

The quantitative importance of these effects is assessed through numerical simulation of the full stochastic model with long-lived households and aggregate productivity risk. I solve the model using finance-informed neural networks (FINNs) and calibrate household preferences, firm technology, and matching parameters to realistic values. Simulating the economy for extended horizons allows me to compute ergodic distributions of all endogenous variables and measure the long-run effects of employer matching on capital accumulation, output, wages, and welfare. The numerical results confirm and quantify the analytical predictions: introducing employer matching at empirically realistic levels reduces equilibrium

equity returns by 79 basis points, increases the aggregate capital stock by 6.1%, and raises wages by 1.7%. I also conduct comparative statics over the matching rate  $\psi$ , demonstrating that the effects scale with the generosity of the matching policy. These quantitative results demonstrate that employer matching has first-order macroeconomic effects that extend well beyond simple redistribution between shareholders and workers.

## 2.2 Related Literature

This paper contributes to several distinct but interconnected literatures spanning macroeconomics, finance, and household portfolio choice. The central methodological contribution—pairing stochastic overlapping generations with neoclassical firm investment to study how retirement policy affects corporate capital accumulation through endogenous equity pricing—bridges literatures that have traditionally developed separately.

### 2.2.1 Overlapping Generations Models and Retirement Policy

The foundation of this work rests on the long tradition of overlapping generations models pioneered by Samuelson (1958). The existence and characterization of Markov equilibria in stochastic OLG models has been studied extensively by Spear and Srivastava (1986), Duffie et al. (1994), and Citanna and Siconolfi (2010), providing the theoretical underpinnings for the equilibrium concept employed in this paper.

Within the OLG framework, a substantial literature examines how retirement policies affect aggregate savings and capital accumulation. Gomes and Michaelides (2003) develop a general equilibrium life-cycle model with incomplete markets and heterogeneous agents to evaluate DB versus DC pension systems, finding that social welfare is maximized at small positive DB levels due to intergenerational risk-sharing. Coimbra et al. (2023) show that the historical shift from DB to DC systems reduced the equity risk premium—not dissimilar to my result that employer matching lowers equity returns (though my model does not have a risk-free asset). Krueger and Kubler (2006) study Pareto-improving social security reforms when financial markets are incomplete, demonstrating that retirement policies can have first-order welfare consequences through general equilibrium channels. While I abstract from social security to isolate the employer matching mechanism, the question of how retirement subsidies affect aggregate savings and capital accumulation remains central to both analyses. Krueger and Kubler (2002) analyze intergenerational risk-sharing via social security in incomplete markets, highlighting the asset pricing im-

plications of retirement policy—a theme that carries through to the employer matching context studied here. Hosseini and Shourideh (2019) develop a quantitative framework to study optimal retirement financing reforms. Their emphasis on the general equilibrium interactions between retirement policy, capital accumulation, and welfare aligns closely with the mechanism I study, though they focus on social security design rather than employer matching.

More recently, A. Zhang (2022) studies the general equilibrium implications of target-date funds, showing that innovations in retirement finance can affect real corporate behavior through general equilibrium price effects. Their finding uses the restrictive portfolio shares of target-date funds as the mechanism for rebalancing and repricing in equilibrium, similar in spirit to the mechanism I identify with employer retirement matching.

The microfoundations of firms’ decisions to sponsor retirement plans—including tax incentives, administrative costs, and labor market competition—have been studied by Bloomfield et al. (2025), who document that tax credits have limited take-up among small firms and identify barriers to retirement plan adoption. I take these microfoundations as given, focusing instead on the general equilibrium consequences of matching once it is in place. This modeling choice allows me to isolate the core mechanism linking retirement policy to firm investment while abstracting from the complex incentives governing plan adoption. Future work could usefully integrate both margins by endogenizing firms’ matching decisions through a bargaining or competition framework.

### 2.2.2 Household Portfolio Choice and Retirement Savings Behavior

The household side of the model builds on the extensive literature examining optimal portfolio choice over the life cycle. Duarte et al. (2021) develop a machine-learning approach to solve for optimal portfolio allocation across stocks, bonds, and liquid accounts in a realistic life-cycle model, finding substantial welfare losses from using simple age-based target-date fund rules rather than customizing to individual circumstances. The literature has long recognized that many households undersave for retirement relative to normative benchmarks, motivating policies to encourage greater retirement contributions. Bhargava and

Conell-Price (2022) conduct a field experiment among 401(k) participants and document present bias in retirement savings decisions, while Laibson (1996) shows theoretically how hyperbolic discounting generates undersaving. Campbell (2016) argues that behavioral biases in financial decision-making justify paternalistic interventions like automatic enrollment and employer matching. In this context, employer matching can be understood as a policy response to undersaving: by subsidizing contributions, matching provides an extra incentive to overcome present bias. My model abstracts from behavioral frictions by assuming time-consistent preferences, allowing me to isolate the pure general equilibrium pricing effect of matching from behavioral responses. The finding that matching generates substantial capital accumulation and welfare gains even in the absence of behavioral frictions suggests that the mechanism operates independently of whether households are initially undersaving.

Dammon, Spatt, and H. H. Zhang (2004) study optimal asset location and allocation when investors face differential tax treatment across taxable and tax-deferred accounts, demonstrating that tax considerations generate substantial portfolio distortions. While I abstract from explicit taxation to maintain tractability, the economic intuition of my model would remain in a more realistic taxation environment. Gomes, Michaelides, and Polkovnichenko (2009) also explores optimal asset location under taxable and tax-deferred retirement accounts, focusing largely on direct versus indirect ownership of stocks.

### 2.2.3 Asset Pricing in Overlapping Generations Economies

A central contribution of this paper is demonstrating how employer matching affects equilibrium asset prices through the SDF. The relationship between demographic structure, household portfolio choice, and asset prices in OLG models has been studied extensively. Constantinides, Donaldson, and Mehra (2002) show that borrowing constraints on young households can help resolve the equity premium puzzle by reducing equity demand from high-marginal-utility young agents, thereby increasing equilibrium risk premia. Storesletten, Telmer, and Yaron (2007) extend this analysis to incorporate idiosyncratic income risk, demonstrating that household heterogeneity in risk exposure affects equilibrium asset prices and quantities.

Gârleanu and Panageas (2015) analyze how heterogeneity in risk aversion and age affects portfolio composition and equilibrium returns, coining the phrase ‘young, old, conservative, and bold’ to describe how life-cycle patterns in risk tolerance shape the cross-sectional distribution of equity holdings. Geanakoplos, Magill, and Quinzii (2004) study how demographic shifts affect long-run stock market returns through cohort-specific demand for equity, finding that predictable changes in the age distribution generate slow-moving variation in equilibrium prices.

In incomplete markets settings like the one I study, the SDF is not unique, raising the question of which discount rate firms should use when evaluating investment projects. Hansen and Jagannathan (1991) develop the concept of an efficient SDF that minimizes the second moment subject to satisfying the pricing restriction, providing a disciplined approach to aggregating heterogeneous household marginal rates of substitution. I adopt their methodology to construct the firm’s discount rate, ensuring computational stability while accurately reflecting household pricing.

#### 2.2.4 Neoclassical Firm Investment and the Cost of Capital

On the firm side, the model builds on the q-theory of investment developed by Tobin (1969) and Jorgenson (1963), with the formal dynamic investment problem following Hayashi (1982).

The integration of households following the overlapping generations structure with firm dynamic investment decisions through the cost of capital is less common in the literature. Most OLG models either take firm investment as exogenous or study it in simplified static settings, and many neoclassical investment models of the firm take the firm’s discount rate as an exogenously set parameter. By pairing a sixty-period OLG model with a neoclassical investment problem and solving for the joint equilibrium using neural networks, I bridge the household and firm sides of the economy in a quantitatively realistic framework. This integration is essential for capturing the general equilibrium mechanism: employer matching affects household equity demand, which alters equilibrium prices, which changes the SDF used by firms, which ultimately affects investment and real outcomes.

### 2.2.5 Computational Methods: Deep Learning for Economics

The computational approach employed in this paper—finance-informed neural networks (FINNs)—builds on recent advances in applying deep learning to solve economic models. Raissi, Perdikaris, and Karniadakis (2019) pioneer physics-informed neural networks (PINNs), which embed partial differential equations directly into the loss function, allowing neural networks to approximate solutions to complex PDEs without requiring large datasets. Azinovic, Gaegauf, and Scheidegger (2019) adapt this approach to economics, developing ‘deep equilibrium nets’ that solve for global policy functions in dynamic stochastic general equilibrium models by embedding equilibrium conditions into the training objective. Azinovic and Zemlicka (2024), whose methodology I follow closely, apply FINNs to stochastic overlapping generations models and demonstrate that the approach can handle high-dimensional state spaces and non-convexities that traditional projection methods struggle with.

More broadly, Fernández-Villaverde (2025) provides a comprehensive introduction to deep learning methods for solving economic models, emphasizing that neural networks offer a flexible alternative to traditional value function iteration, perturbation, and projection methods. L. Maliar, S. Maliar, and Winant (2021) compare deep learning approaches to other solution methods across a range of macroeconomic models, finding that neural networks perform particularly well in high-dimensional settings with nonlinearities—precisely the environment studied in this paper. Han, Yang, and E (2021) develop specialized methods for summarizing rich economic information in heterogeneous agent models with aggregate shocks by endogenously identifying key moments of the model, providing further evidence that deep learning can tractably handle the computational challenges arising in modern quantitative macroeconomics.

## 2.3 The Model

I pair an overlapping generations model with one representative type per generation on the household side with a neoclassical model of firm investment on the production side. There are  $I$  periods of life for each generation. Lifetimes are deterministic.

Available for purchase at all stages of life are financial assets for the purpose of consumption smoothing and retirement financing, taking the form of equity shares in the representative firm. Discussion of these assets is also included below. Households inelastically supply labor to the firm in exchange for competitively determined wages for labor.

The firm has already had an (unmodeled) initial equity issuance with the number of shares normalized to 1. Households trade fractional shares at competitively determined prices, and the shares entitle the households to one-period forward dividends.

For any variable  $x$ , subscripts  $x_{i,t}$  means the value of  $x$  for the  $i$ -th oldest generation; i.e. the generation born at  $t - i$ . Additionally, for any lower-case  $i, t$ -subscripted variable, let the upper-case script variable denote its  $i$ -vector:  $\mathcal{X}_t := \{x_{i,t}\}_i$ .

### 2.3.1 Market Structure

The source of aggregate risk in the economy will be the total factor productivity affecting firm output, an autoregressive process of degree one –  $AR(1)$  – in logs. Log-TFP follows the following process:

$$\log Z_{t+1} = \rho_Z \log Z_t + \sigma_Z \epsilon_{t+1} \quad (2.2)$$

where  $0 < \rho_Z < 1$  so that  $\log Z$  has a stationary distribution and the shock  $\epsilon$  is drawn from a standard Normal Distribution:  $\epsilon_Z \sim \mathcal{N}(0, 1)$ . Capital depreciation will also depend on the aggregate shock. In particular, the stochastic depreciation process will be inversely correlated with the TFP process and will

be parameterized by  $\xi_\delta$ .

$$\delta(Z_t) = \bar{\delta} \times \frac{2}{2 + \xi_\delta(Z_t - 1)} \quad (2.3)$$

As discussed above, I model one asset. Households trade shares of equity in the firm which are normalized to one share. Short sales on equity are forbidden, so that  $s_{i,t} \geq 0$  for all  $i, t$ .

The gross return on equity is given as  $R_t = \frac{p_t + d_t}{p_{t-1}}$  where  $p_t$  is the competitively determined time- $t$  price of the equity shares and  $d_t$  is the time- $t$  dividend of the firm. More discussion of the firm will take place below.

### 2.3.2 Households

As mentioned above, each generation lives for exactly  $I$  periods. Using  $t$  as the time index, the index  $i \in \{0, \dots, I - 1\}$  represents age concurrently with time.

Households are endowed with labor efficiency units  $\ell_i$  over their life-cycle for all  $i, t$ .

In each period, households choose optimal consumption and savings in risky equity, subject to the constraints as described in the previous section.

Households earn income from wages:

$$y_{i,t} = w_t \ell_i \quad (2.4)$$

at wage rate  $w_t$  per efficiency unit of labor. Additionally, firms contribute retirement savings to households at a known schedule, matching some fraction  $\psi \geq 0$  of contribution up to a threshold given as a percent of income  $\phi$ .

The sequential budget constraints for households alive at time  $t$  are thus:

$$c_{i,t} + p_t s_{i,t} = y_{i,t} + (p_t + d_t) a_{i,t} \quad (2.5)$$

$$a_{i,t} = s_{i-1,t-1} + m_{i-1,t-1} \quad (2.6)$$

$$m_{i,t} = \psi \min \left\{ s_{i,t}, \frac{\phi y_{i,t}}{p_t} \right\} \quad (2.7)$$

Since households die with certainty after their last period of life and are endowed with no financial assets:

$$s_0 = s_{I-1} = 0 \text{ for all } t.$$

Households form utility over lifetime consumption:

$$U : \mathbb{R}_{++}^I \rightarrow \mathbb{R}$$

I'll assume additively time-separable von Neumann–Morgenstern expected constant relative risk aversion (CRRA) utility, with coefficient of relative risk aversion  $\gamma$ :

$$U(\{c_{i,t+i}\}_{i=0}^{I-1}) = \mathbb{E}_t \sum_{i=0}^{I-1} \beta^i u(c_{i,t+i}) : u(c) = \begin{cases} \frac{c^{1-\gamma}-1}{1-\gamma} & \gamma \in \mathbb{R}_+ \setminus \{1\} \\ \log(c) & \gamma = 1 \end{cases}$$

Then the households' formal problem can be written as:

$$\max_{\{c_{i,t+i}, s_{i,t+i}\}_{i=0}^{I-2}} \left\{ U(\{c_{i,t+i}\}_{i=0}^{I-1}) \right\} \text{ s.t.} \quad (2.8)$$

$$c_{i,t} + p_t s_{i,t} = y_{i,t} + (p_t + d_t) \left( s_{i-1,t-1} + \psi \min \left\{ s_{i,t}, \frac{\phi y_{i,t}}{p_t} \right\} \right) \quad (2.9)$$

$$s_{i,t} \geq 0 \quad (2.10)$$

Let  $\mu$  be the Lagrange multiplier on the short selling constraint. Households' optimality conditions are given by the following Euler equations and KKT conditions:

$$1 = \mathbb{E}_t \left[ \frac{\beta u'(c_{i+1,t+1})}{u'(c_{i,t})} \cdot \frac{p_{t+1} + d_{t+1}}{p_t} \cdot \left( 1 + \frac{\partial m_{i,t}}{\partial s_{i,t}} \right) \right] + \mu_{i,t} \quad (2.11)$$

$$s_{i,t} \geq 0 \quad (2.12)$$

$$\mu_{i,t} s_{i,t} = 0 \quad (2.13)$$

### 2.3.3 Firms

The economy is populated by a representative firm. The firm hires a measure  $L$  of workers and makes a dynamic investment decision  $I_t$ , accumulating its own stock of capital  $K_t$  depreciating at rate  $\delta_t$  to maximize the expected discounted payout of dividends  $d_t$  at discount rate  $\Lambda_t$ . The firm pays a convex adjustment cost  $\Phi(I_t, K_t)$  and contributes aggregate matches  $M_t$  to employee retirement funds. The firm takes the equity price  $p_t$  and the matches  $m_{i,t}$  as given and does *not* internalize the effect of matching on repricing through the SDF. The firm's period budget constraint is thus:

$$d_t = Y_t - I_t - w_t L_t - p_t M_t - \Phi(I_t, K_t) \quad (2.14)$$

$$Y_t = Z_t F(K_t, L_t) \quad (2.15)$$

$$I_t = K_{t+1} - (1 - \delta_t) K_t \quad (2.16)$$

$$\Phi(I_t, K_t) = \frac{\eta}{2} \left( \frac{I_t}{K_t} - \bar{\delta} \right)^2 K_t \quad (2.17)$$

$$M_t = \sum_{i=0}^{I-1} m_{i,t} \quad (2.18)$$

The firm's production technology is Cobb-Douglas with parameter  $\alpha$  and Total Factor Productivity multiplier  $\bar{Z}$ :

$$F(K_t, L_t) = \bar{Z} K_t^\alpha L_t^{1-\alpha} \quad (2.19)$$

and then total output is given as  $Y_t = Z_t F(K_t, L_t)$ .

The firm dividends  $d_t$  are paid to the equity holders every period. The firm discounts its dividends at the stochastic discount rate of the equity holders,  $\Lambda_{t+1}$ . Discussion of its calculation will follow below.

Define the cumulative SDF from  $s$  to  $t$  as

$$\Lambda_{s:t} = \prod_{\tau=s+1}^t \Lambda_\tau \quad (2.20)$$

The formal statement of the firm's problem is then given as:

$$\max_{\{L_t, I_t, K_{t+1}\}_{t=0}^{\infty}} \left\{ \mathbb{E}_0 \left[ \sum_{t=0}^{\infty} \Lambda_{0:t} d_t \right] \right\} \quad \text{s.t.} \quad (2.21)$$

$$d_t = Z_t F(K_t, L_t) - I_t - w_t L_t - p_t M_t - \Phi(I_t, K_t) \quad (2.22)$$

$$I_t = K_{t+1} - (1 - \delta_t) K_t \quad (2.23)$$

Let  $Q_t$  be the Lagrange multiplier on the law of motion for capital. The solution to the firm's problem is the following set of Euler equations:

$$Q_t = \mathbb{E}_t \left[ \Lambda_{t+1} (Z_{t+1} F_K(K_{t+1}, L_{t+1}) + (1 - \delta_{t+1}) Q_{t+1} - \Phi_K(I_{t+1}, K_{t+1})) \right] \quad (2.24)$$

$$Q_t = 1 + \Phi_I(I_t, K_t) \quad (2.25)$$

The firm labor decision results in a static problem:

$$Z_t F_L(K_t, L_t) = w_t \quad (2.26)$$

### 2.3.4 Market Clearing

The equity and labor markets need to clear:

$$\sum_i e_{i,t} = 1 \qquad \sum_i \ell_{i,j} = L_t = 1 \quad (2.27)$$

The consumption market will clear via Walras's Law.

### 2.3.5 Equilibrium

The equilibrium of study in long-lived stochastic overlapping generations models is that of a recursive Markov equilibrium.

**Definition 2.** *The **recursive Markov equilibrium** is defined by time-homogeneous policy functions for the purchase and pricing of assets for state variable  $\chi$ :*

$$\{S(\chi), I(\chi), P(\chi), \mu(\chi)\}$$

where  $\chi \supseteq (\mathcal{A}, \mathcal{K}, Z)$  is taken to be at a minimum the lagged asset holdings of all households and the current realization of the shock. The time-homogeneous policy functions solve the households' problems and the firm's problem. Markets for capital, bonds, and labor must clear. Feasibility arises from Walras's Law and is given by

$$Y_t = C_t + I_t + \Phi_t$$

where  $C_t$  is aggregate consumption,  $I_t$  is aggregate investment, and  $\Phi_t$  is the adjustment cost paid by the firm in period  $t$ .

From equilibrium relationships, a few restrictions can be made immediately that will assist in computing the full model. For instance, the wage can be solved as a static problem since the labor supply is inelastic, so the market-clearing wage is given as:

$$w_t = (1 - \alpha) \frac{Y_t}{L_t} \tag{2.28}$$

where  $L = \sum_i \ell_i = 1$ .

## 2.4 Two-Period Steady State Equilibrium Analysis

The model's steady state can be solved in closed form without aggregate risk and with two generational cohorts: young and old.

### 2.4.1 Households

A unit mass of identical households lives for two periods. The young household inelastically supplies a unit mass of labor for competitively determined wages  $w$ . Households choose savings  $s \geq 0$  and receive a (capped) savings match:  $m(s) = \psi \min\left\{s, \phi \frac{w}{p}\right\}$  where  $\psi \geq 0$  is the match rate parameter and  $\phi$  is the match cap parameter measured as a fraction of wage income valued at the asset price  $p$ . The household budget constraints are given as:

$$c_y = w - ps \quad (2.29)$$

$$c_o = (p + d)(s + m(s)) \quad (2.30)$$

and preferences are logarithmic:

$$U(s) = \log c_y + \beta \log c_o \quad (2.31)$$

The formal household problem is given as

$$\max_{c_y, c_o, s} \{U(s)\} \text{ subject to budget constraints} \quad (2.32)$$

There are two cases: when  $s < \frac{\phi w}{p}$  will be referred to as *Case 1* and when  $s > \frac{\phi w}{p}$  will be *Case 2*.

*Case 1.* When  $s < \frac{\phi w}{p}$ , the Euler equation is given as:

$$\frac{p}{c_y} = \beta \frac{p + d}{c_o} (1 + \psi) \implies s^* = \frac{\beta w}{(1 + \beta)p} \quad (2.33)$$

Reconciling the bounds, this holds when  $\beta < \frac{\phi}{1 - \phi}$ .

*Case 2.* When  $s > \frac{\phi w}{p}$ , the Euler equation is given as:

$$\frac{p}{c_y} = \beta \frac{p + d}{c_o} \implies s^* = \frac{(\beta - \psi\phi)w}{(1 + \beta)p} \quad (2.34)$$

Reconciling the bounds, this holds when  $\beta > \frac{\phi(1+\psi)}{1-\phi}$ .

*Case 3.* At the kink,  $s = \frac{\phi w}{p}$ . This occurs when  $\frac{\phi}{1-\phi} < \beta < \frac{\phi(1+\psi)}{1-\phi}$ .

We can summarize the household demand schedule as:

$$s^*(p) = \begin{cases} \frac{\beta w}{(1 + \beta)p}, & \beta \leq \frac{\phi}{1-\phi} \\ \frac{\phi w}{p}, & \frac{\phi}{1-\phi} < \beta < \frac{\phi(1+\psi)}{1-\phi} \\ \frac{w(\beta - \psi\phi)}{(1 + \beta)p}, & \beta \geq \frac{\phi(1+\psi)}{1-\phi} \end{cases} \quad (2.35)$$

In equilibrium, the asset market must clear:  $s + m(s) = 1$ . Using the household demand schedule and solving for the price  $p$  in each regime gives

$$p^* = \begin{cases} \frac{\beta(1 + \psi)}{1 + \beta} w, & \beta \leq \frac{\phi}{1-\phi} \\ (1 + \psi) \phi w, & \frac{\phi}{1-\phi} < \beta < \frac{\phi(1+\psi)}{1-\phi} \\ \frac{\beta(1 + \psi\phi)}{1 + \beta} w, & \beta \geq \frac{\phi(1+\psi)}{1-\phi} \end{cases} \quad (2.36)$$

Finally, substituting the market-clearing prices back into the demand schedules provides the equilibrium allocations:

$$(s^*, m^*) = \begin{cases} \left( \frac{1}{1 + \psi}, \frac{\psi}{1 + \psi} \right), & \beta \leq \frac{\phi(1+\psi)}{1-\phi} \\ \left( \frac{\beta - \psi\phi}{\beta(1 + \psi\phi)}, \frac{\psi\phi(1 + \beta)}{\beta(1 + \psi\phi)} \right), & \beta \geq \frac{\phi(1+\psi)}{1-\phi} \end{cases} \quad (2.37)$$

### 2.4.2 Representative Firm

A representative firm operates a constant-returns-to-scale Cobb-Douglas technology with parameter  $\alpha \in (0, 1)$ :

$$Y = ZK^\alpha L^{1-\alpha} \quad (2.38)$$

It hires labor  $L$  competitively at wage  $w$  and owns its capital stock  $K$ . Assume no depreciation:  $\delta = 0$  so that investment is  $I = K' - K$ . Dividends are profits net of matches:

$$d = Y - wL - pm - I \quad (2.39)$$

where  $m \in [0, 1]$  is the per-share matched quantity determined by households and funded by the firm at market value  $p$ . The firm chooses labor quantity  $L$  and (forward) capital  $K'$  to maximize the discounted flow of dividends to the owners of the firm. Since this is a model without aggregate risk, the discount rate of the owners of the firm is the return received by households, the *cost of equity*:

$$R = \frac{p + d}{p} \quad (2.40)$$

The formal statement of the Firm's Problem is recursively given as:

$$V(K) = \max_{L, K'} \left\{ d + \frac{1}{R} V'(K') \right\} \quad (2.41)$$

First, the wage is competitively determined as

$$w = (1 - \alpha)ZK^\alpha \quad (2.42)$$

when  $L = 1$ . And the firm faces capital Euler equation given by

$$R = 1 + \alpha Z(K')^{\alpha-1} \quad (2.43)$$

We can now substitute equilibrium quantities for  $wL = (1 - \alpha)ZK^\alpha$  back into the firm's dividend and use this to simplify the cost of equity:

$$R = 1 + \frac{\alpha Z K^\alpha}{p} - m \quad (2.44)$$

Now moving to steady state so that  $K = K'$  and reconciling equations (2.43) and (2.44):

$$1 + \frac{\alpha Z K^\alpha}{p} - m = 1 + \alpha Z K^{\alpha-1} \quad (2.45)$$

Recall equation (2.36) from the solution to the household problem and substitute in that  $w = (1 - \alpha)ZK^\alpha$ :

$$p^* = \begin{cases} \frac{\beta(1+\psi)}{1+\beta}(1-\alpha)ZK^\alpha, & \beta \leq \frac{\phi}{1-\phi} \\ (1+\psi)\phi(1-\alpha)ZK^\alpha, & \frac{\phi}{1-\phi} < \beta < \frac{\phi(1+\psi)}{1-\phi} \\ \frac{\beta(1+\psi\phi)}{1+\beta}(1-\alpha)ZK^\alpha, & \beta \geq \frac{\phi(1+\psi)}{1-\phi} \end{cases} \quad (2.46)$$

Notice that in all cases, we can express  $p^* = c(\psi)K^\alpha$  where  $c(\psi)$  represents the coefficients and  $c$  is strictly increasing in  $\psi$ :  $c'(\psi) > 0$ . Now substituting  $p^* = c(\psi)K^\alpha$  into (2.45) yields:

$$\frac{\alpha Z K^\alpha}{c(\psi)K^\alpha} - m = \alpha Z K^{\alpha-1} \quad (2.47)$$

implying

$$K^* = \left( \frac{\alpha Z c(\psi)}{\alpha Z - m c(\psi)} \right)^{\frac{1}{1-\alpha}} \quad (2.48)$$

**Theorem 2.4.1** (Matching Raises the Capital Stock).  $\frac{\partial K^*}{\partial \psi} > 0$ .

*Proof.* Define  $x = K^{1-\alpha} = \frac{\alpha Z}{\frac{\alpha Z}{c(\psi)} - m}$ . Then compute:

$$\frac{\partial x}{\partial \psi} = -\frac{\alpha Z}{\left( \frac{\alpha Z}{c(\psi)} - m \right)^2} \times \left( -\frac{\alpha Z}{c(\psi)^2} c'(\psi) - m'(\psi) \right) > 0 \quad (2.49)$$

Notice that the negative signs cancel out, the denominator of the first time must be positive by positivity of  $K^*$ , and both  $c'$  and  $m'$  are positive. Since  $\frac{\partial K}{\partial x} > 0$ , we have proven the claim.  $\square$

**Theorem 2.4.2** (Matching Lowers the Equilibrium Returns).  $\frac{\partial R^*}{\partial \psi} < 0$ .

*Proof.* Recall equation (2.44) and substitute in the price as before:

$$R = 1 + \frac{\alpha Z K^\alpha}{p} - m \quad (2.50)$$

$$= 1 + \frac{\alpha Z K^\alpha}{c(\psi) K^\alpha} - m \quad (2.51)$$

$$= 1 + \frac{\alpha Z}{c(\psi)} - m \quad (2.52)$$

Taking derivatives immediately yields

$$\frac{\partial R}{\partial \psi} = -\frac{\alpha Z}{c(\psi)^2} c'(\psi) - m'(\psi) < 0 \quad (2.53)$$

Where the same logic as the proof of  $\frac{\partial R^*}{\partial \psi} < 0$  holds.  $\square$

### 2.4.3 Marginal vs. Average Tobin's Q

Starting from (2.45) and collecting like terms:

$$\alpha Z K^{\alpha-1} \left( \frac{K}{p} - 1 \right) = m \quad (2.54)$$

Notice immediately that in the case when  $\psi = 0 \implies m = 0$  and employers do not match, the solution to equation (2.54) is  $K^* = p$ , the celebrated Hayashi 1982 result. However, with employer matching we have  $m > 0$  and therefore  $K^* > p$ . The employer matching therefore introduces a wedge between average and marginal  $Q$ , something that may be of note to empiricists studying corporate finance.

#### 2.4.4 Worker-Financed Match

A potential critique of the baseline model is that the shareholder-financing assumption—whereby employer matching is paid out of dividends—may mechanically drive the reduction in equilibrium returns through decreased dividends. To address this concern, I consider an alternative specification in which workers finance the employer match through reduced wages rather than shareholders financing it through reduced dividends. Specifically, I impose the constraint  $(1 - \alpha)Y = wL + pm$ , which forces total labor compensation (wages plus matching contributions) to equal labor’s marginal product in a match-free policy environment, effectively making workers bear the incidence of the matching cost.

The mathematical derivations that follow demonstrate that even under worker financing, the central mechanism of the paper remains intact: employer matching still increases the SDF, reduces equilibrium returns, and stimulates firm capital investment. The key economic insight—that matching affects capital accumulation through endogenous equity pricing and the SDF—is therefore robust to the financing assumption.

The reason I do not adopt worker financing or shared financing as the primary specification is purely computational. Under shareholder financing, the wage can be solved statically as  $w_t = (1 - \alpha)Y_t/L_t$  from the firm’s first-order condition for labor, which greatly simplifies the numerical solution. Under worker financing, the wage becomes a more complex *policy function* that must satisfy the constraint  $(1 - \alpha)Y = wL + pm$  and therefore depends on the equilibrium matching contribution and equity price. This requires solving for the wage as part of the policy function rather than as a closed-form static relationship, increasing the computational burden when using neural network approximations. Since the economic mechanism is unchanged, I prioritize computational efficiency by maintaining shareholder financing in the main analysis while demonstrating robustness to alternative incidence assumptions below. While this modeling choice may affect the quantitative interpretation of the numerical results, it is worth noting that any quantitative calibration in a tax-free environment necessarily abstracts from first-order features of actual retirement policy, limiting the precision of welfare calculations regardless of the inci-

dence assumption. More broadly, microfounding the choice of matching parameters and the source of financing through explicit bargaining models or contract theory would constitute a rich area for future research, allowing labor economists and bargaining theorists to build on the general equilibrium framework developed here to endogenize retirement policy design.

To show that the source of financing does not impact the main result, assume that workers fully finance their retirement match:  $wL + pm = (1 - \alpha)Y$ . The firm's dividend then simplifies to  $d = Y - wL - pm = \alpha ZK^\alpha$  after steady state is imposed and equilibrium conditions are substituted in. The return then becomes:

$$R = 1 + \frac{\alpha ZK^\alpha}{p} \quad (2.55)$$

while the firm's Euler equation is unchanged:

$$R = 1 + \alpha ZK^{\alpha-1} \quad (2.56)$$

As before, note that we can write  $p = c(\psi)w$ . Substituting in the (new) equilibrium restriction that  $w = Z(1 - \alpha)K^\alpha - pm$  and then grouping terms for  $p$  yields the equilibrium price:

$$p^* = \frac{c(\psi)Z(1 - \alpha)}{1 + c(\psi)m(\psi)} K^\alpha \quad (2.57)$$

Reconciling equations (2.55) and (2.56) and substituting in (2.57) yields the expression for equilibrium levels of capital:

$$K^* = \left( Z(1 - \alpha) \frac{c(\psi)}{1 + m(\psi)c(\psi)} \right)^{\frac{1}{1-\alpha}} \quad (2.58)$$

Carefully accounting for the parameter regions that dictate the values  $c(\psi)$  and  $m(\psi)$ , one can show that this expression is increasing in  $\psi$  in equilibrium. Finally, since  $K^*$  is increasing in  $\psi$  it is easy to show that  $R$  is decreasing in  $\psi$ , which recovers the result from the dividend-financed matching baseline model. The formal proof is tedious but can be verified by the reader.

## 2.5 Sixty-Period Computation

The previous section showed the central mechanism of this paper in closed form, while this section will quantify the mechanism in a more quantitatively realistic model. The main model objects that need external calibration from the data are the life-cycle profiles of labor efficiency units. Further details on those below. Other model parameters are calibrated as follows:

Table 2.1: Calibrated Parameters

| Description                          | Symbol         | Value            | Target / Source         |
|--------------------------------------|----------------|------------------|-------------------------|
| Periods of life                      | $I$            | 60               | –                       |
| Efficiency units of labor            | $\ell_i$       | See Figure (2.1) | SCF (2019)              |
| Discount factor                      | $\beta$        | 0.925            | –                       |
| Relative risk aversion               | $\gamma$       | 1                | –                       |
| Capital share (Cobb–Douglas)         | $\alpha$       | 0.30             | –                       |
| TFP multiplier                       | $\bar{Z}$      | $10^{1-\alpha}$  | Normalization           |
| Adjustment Cost                      | $\eta$         | 2                | –                       |
| Persistence of log TFP               | $\rho_Z$       | 0.90             | –                       |
| Std. dev. of TFP shock               | $\sigma_Z$     | 0.0087           | $\text{std}(Z_t) = 2\%$ |
| Average depreciation rate            | $\bar{\delta}$ | 0.10             | –                       |
| Depreciation sensitivity (agg. risk) | $\xi_\delta$   | 2                | –                       |
| Matching Rate                        | $\psi$         | $\{0, 0.5, 1\}$  | –                       |
| Matching Cap Rate                    | $\phi$         | 0.06             | –                       |

It is worth noting a few of the calibrated parameter values.

- $\beta = 0.925$  is chosen for two reasons: computational stability and to reflect that one of the stated goals of employer retirement matching is to incentivize (otherwise impatient) households to save for retirement earlier. With respect to computational stability, the model’s link between household’s SDF and the firm’s real decisions loses stability as the households discount by less.
- $\psi = 0, 0.5, 1$  reflect the following matching offers, common among large employers:
  - no matching,
  - ‘a match of 50¢ on the dollar, up to 6% of pre-tax income,’ and

- ‘a dollar-for-dollar match, up to 6% of pre-tax income’
- $\xi_\delta = 2$  is based on Azinovic and Zemlicka (2024) and troubleshooting to stabilize the firm’s investment problem with counter-cyclical stochastic depreciation.
- $\gamma = 1$  (logarithmic utility, implying unit intertemporal elasticity of substitution) serves both computational and economic purposes. Computationally, it aids numerical stability by avoiding extreme curvature in household marginal utilities. Economically, it isolates the general equilibrium pricing effect of employer matching from confounding intertemporal substitution responses: with unit IES, variations in savings behavior across matching regimes arise purely from the matching subsidy term in the SDF rather than from households adjusting their willingness to substitute consumption intertemporally in response to changing returns. This parsimonious specification cleanly demonstrates that the matching mechanism operates through endogenous equity pricing rather than through household elasticity channels.

### 2.5.1 Data Sources and Calibration

This section describes my approach to estimating empirically realistic life-cycle profiles of labor efficiency units from micro-data and explains the methodological choices underlying the calibration.

I use the Survey of Consumer Finances (SCF) 2019 release (Board of Governors of the Federal Reserve System 2019) to construct age-specific labor income profiles, which I use to calibrate the labor efficiency units  $\{\ell_i\}_i$  in the model. I use the 2019 release rather than more recent data because it reflects labor market conditions prior to the COVID-19 pandemic, providing a more stable baseline for calibrating steady-state income profiles.

A well-known feature of the SCF is that it oversamples high-wealth and high-income households to improve estimates of top-tail distributions. While this design is valuable for studying wealth concentration, it poses challenges for calibrating a representative-agent model intended to capture the behavior of typical workers. To address this, I truncate the top 15% of the income distribution before constructing the

age-income profile. This truncation serves two purposes. First, it mitigates the mechanical bias introduced by oversampling: without truncation, the mean income profile would be upward-biased relative to the population average, distorting the calibration of labor efficiency units. Second, and more importantly for the economic mechanism of this paper, ultra-wealthy households are unlikely to face binding employer matching constraints. Employer matching contributions are typically capped at a fraction of pre-tax income—commonly 6% in practice, reflected in my calibration of  $\phi = 0.06$ . For households with very high incomes, this cap implies that the absolute dollar value of foregone matching is trivial relative to total wealth, making it economically implausible that such households would adjust their retirement savings behavior in response to matching incentives. Put differently, the matching mechanism I study operates primarily through middle- and upper-middle-income workers who face meaningful tradeoffs between current consumption and capturing the employer match. By excluding the top 15%, I focus the calibration on the income range where matching incentives are most likely to bind and affect savings behavior. Future work could usefully study heterogeneity in employer matching generosity across income groups and model transitions between income classes and their associated retirement plans, which would allow for richer analysis of how matching affects wealth inequality and mobility over the life cycle.

Figure 2.1 displays the resulting age-income profile. The profile exhibits the characteristic hump shape documented extensively in labor economics: income rises steeply from ages 20-40 as workers accumulate human capital and climb career ladders, peaks in the mid-40s to early-50s, and declines gradually as workers approach retirement. The hump shape has direct implications for retirement matching: households face the highest incentive to match when young, since low incomes means otherwise low savings. In the model implementation, I normalize these income values so that the labor efficiency units satisfy  $\sum_i \ell_i = 1$ .

Since samples are relatively small when segmented by age, the raw age-income profile is not smooth. Without smoothing, the resulting life-cycle profile exhibits substantial noise that complicates numerical solution of the model. Neural networks trained on noisy input data may have difficulty isolating the equilibrium policy functions from sampling variation, potentially reducing the accuracy of the com-

puted equilibrium. To address this, I smooth the raw income profile using a cubic polynomial before feeding it into the model. The cubic specification is parsimonious yet flexible enough to capture the well-established hump shape: the quadratic term allows the profile to rise and then fall, while the cubic term provides additional curvature to fit the asymmetric shape of the hump (steep rise, gradual decline). Higher-order polynomials would risk overfitting the noise inherent in the SCF sample, while lower-order specifications (e.g., quadratic) fail to capture the asymmetry of the empirical profile. As shown in Figure 2.1, the cubic fit provides a smooth approximation that faithfully represents the hump-shaped pattern while eliminating high-frequency variation.




Figure 2.1: Age-efficiency profile calibrated to SCF 2019. Values normalized so that  $\sum_i \ell_i = 1$ .

### 2.5.2 Solution Method: FINNs

The general structure of the algorithm used to solve the fully stochastic model is inspired by Azinovic, Gaegauf, and Scheidegger 2019 and Azinovic and Zemlicka 2024, though with a few differences detailed below. For the most up-to-date primer on deep learning for economics, including first-principles building of (deep) neural networks, see Fernández-Villaverde 2025.

Following the physics-informed neural networks (PINNs) approach of Raissi, Perdikaris, and Karniadakis (2019), which embeds partial differential equations directly into the neural network loss function, I develop what I term finance-informed neural networks (FINNs). Rather than the PDEs governing a *physical* system, FINNs incorporate the equations governing the *financial* model—Euler equations, market clearing conditions, and other terms discussed below—as penalty terms in the loss function. This

approach ensures that the trained FINN respects the economic structure of the problem while approximating the equilibrium policy functions.

### 2.5.3 Applications of Domain Knowledge

Naïve machine learning methods excel at finding solutions that adhere to the minimization problem of the loss function as specified by the user. The upside of this means that economists need not spend significant time writing the low-level algorithm needed for training the FINN. The downside is that the landscape may admit low-error pathological ‘solutions’ that evaluate to a very small error but are economically implausible. Examples of this that I encountered in training often include:

- constant or near-constant predictions (usually the capital level or the price of the equity asset) that fail to vary with the underlying shock process in any way;
- predictions that vary only as a monotonic transformation of the underlying shock process (if adding a variance penalty for capital, it refuses to smooth out even in the presence of large adjustment costs);
- and other pathological outcomes.

Especially in the stochastic overlapping generations environment where Euler equation iteration is not necessarily a contraction mapping, these outcomes must be monitored.

To avoid the pathological predictions of the type discussed in the preceding paragraph, I employ the following custom strategies to augment the off-the-shelf routines available in coding software.

1. *Augmenting the state variable with ‘redundant’ information, particularly asset pricing and intertemporal values.* Theory indicates that the state variable required for computation of the model in this paper is  $\mathcal{X}_t = (\{a_{i,t}\}_i, K_t, Z_t)$ . In practice, I found that to break the FINN’s pathological tendency to return predictions that depend only on the shock process itself, including asset pricing values contributed greatly to the network’s ability to converge to economically plausible

equilibria. As such, I model the state variable instead as:

$$\mathcal{X}_t = (\{a_{i,t}\}_i, \mathbb{E}_{t-1}[\Lambda_t], p_{t-1}, Q_{t-1}, K_t, Z_t) \quad (2.59)$$

adding the lagged synthetic risk-free bond price (conditionally expected forward SDF), lagged equity price, and lagged Tobin's Q. Note that this approach lies in stark contrast with the model reduction techniques in the vein of Krusell and Smith 1998: I am actually *expanding* the computational model! Along the same vein, I chose to solve for the following policy functions:

$$\mathcal{Y}_t = (\{s_{i,t}\}_i, p_t, \mathbb{E}_t[\Lambda_{t+1}], I_t) \quad (2.60)$$

In the Q-theory framework, I could equivalently have solved for forward capital, Tobin's Q itself, or the investment rate  $I_t/K_t$  but found that the FINN was most performant when solving for the investment level directly. While not strictly necessary, the synthetic bond price (conditionally expected forward SDF) provides rich information that assists in the simultaneous determination of the firm's investment level.

2. *Using the loss function to discipline economically plausible outcomes.* The formal statement of the loss function is below, but it can be understood as comprising of these parts:

- (a) The mean-square-error (across households and time) of the Fisher-Burmeister transform between the Euler equation residual and period savings;
- (b) The mean-square-error (across time) of the firm's Euler equation;
- (c) The mean-square-error (across time) of the equity market clearing constraint;
- (d) The mean-square-error (across time) of the FINN-predicted synthetic bond price (labeled  $\mathbb{E}_t^{\text{NN}}[\Lambda_{t+1}]$ <sup>1</sup> compared with the Gauss-Hermite quadrature-computed expected value of the

---

<sup>1</sup>Note: this is computed as a scalar output.

true SDF (labeled  $\mathbb{E}_t[\Lambda_{t+1}^*]$ )<sup>2</sup>; and

- (e) The mean-square-error (across time) of the aggregate feasibility constraint. This ensures the pathological solution of  $c_{i,t} \rightarrow 0$  for any  $i, t$  is not found.

3. *Applying the Fisher-Burmeister transformation to eliminate the computation of Lagrange multipliers.* Since the agents in the model are constrained to not sell equity short, the Euler equation may not bind endogenously for all agents. As formulated in the formal model earlier in this paper, this entails solving for Lagrange multipliers of each agent, effectively doubling the required state space. To avoid this, Azinovic and Zemlicka (2024) applies the Fisher-Burmeister transform as follows:

$$\Psi^{\text{FB}}(a, b) = a + b - \sqrt{a^2 + b^2} \quad (2.61)$$

which embeds the KKT conditions as  $a, b \geq 0$  and  $ab = 0 \iff \Psi(a, b) = 0$ . However, since the FINN takes derivatives of equilibrium conditions, the Fisher-Burmeister equation becomes unstable as  $a, b \rightarrow 0$  jointly: the derivative of the square-root approaches positive infinity. To stabilize this, I modify the Fisher-Burmeister equation to:

$$\Psi(a, b; \lambda_{\text{FB}}, \epsilon_{\text{FB}}) = \lambda_{\text{FB}} \left( a + b - \frac{a^2 + b^2}{\sqrt{a^2 + b^2 + \epsilon_{\text{FB}}^2}} \right) + (1 - \lambda_{\text{FB}})a^+b^+ \quad (2.62)$$

where for this model I set  $\lambda_{\text{FB}} = 0.8$  and  $\epsilon_{\text{FB}} = 10^{-3}$ .

4. *Time  $t = 0$  initialization near or on the equilibrium manifold.* Since—as stated previously—iteration on the Euler equations in stochastic overlapping generations models is *not* a contraction mapping, the starting guess is very important to finding equilibria. After each simulation step of drawing the ergodic time series, I will save and store  $\bar{\mathcal{X}}$  as the mean across time of the FINN inputs. Since the FINN is learning the global policy functions, these values can only be expected to evaluate correctly on admissible equilibrium values. In the subsequent time series iteration, I

---

<sup>2</sup>More on the computation of the state-by-state  $\Lambda_{t+1}^*$  below.

will initialize endogenous state variables at time  $t = 0$  based on the values in  $\bar{\mathcal{X}}$ .

5. *Alleviating the sequential bottleneck in data generation for training.* Unlike in many machine learning environments, the data on which the algorithm trains for this model is itself generated by the model. In particular, the policy function predicts forward capital in this model. In order to generate training data, I must simulate outcomes of the aggregate risk process along with predictions for forward capital, which become the capital in the next period. Usual sample sizes for models like this are for around  $T = 10000$  periods. Rather than draw one block of 10000 periods, I instead draw 100 blocks of 100 periods *in parallel*, enormously parallelizing the simulation pass for each loop in the training algorithm.
6. *Smoothing out the hard min function parameterizing the employer match.* Using the parameterization of the employer match as given in the formal model above leads to instability in training: the derivative is discontinuous at the kink point. I use the following smooth approximation::

$$\bar{m}_{i,t} = \phi \frac{w_t \ell_i}{p_t} \quad (2.63)$$

$$m_{i,t} = -\psi \tau \log \left( \exp \left( -\frac{s_{i,t}}{\tau} \right) + \exp \left( -\frac{\bar{m}_{i,t}}{\tau} \right) \right) \quad (2.64)$$

$$\frac{\partial m_{i,t}}{\partial s_{i,t}} = \psi \frac{1}{1 + \exp \left( -\frac{\bar{m}_{i,t} - s_{i,t}}{\tau} \right)} \quad (2.65)$$

As  $\tau \rightarrow 0$ , the matching function approaches the original minimum function and the derivative approaches the step function. My selected value of  $\tau = 10^{-3}$  is small enough to approximate the match closely, and large enough to maintain smoothness.

#### 2.5.4 FINN Architecture and Training

As mentioned above, the state variable required for computation of the model in this paper is  $\mathcal{X}_t = (\{a_{i,t}\}_i, \mathbb{E}_{t-1}[\Lambda_t], p_{t-1}, Q_{t-1}, K_t, Z_t)$ , and the predictions are  $\mathcal{Y}_t = (\{s_{i,t}\}_i, p_t, \mathbb{E}_t[\Lambda_{t+1}], I_t)$ . The FINN architecture consists of two hidden layers, each with 100 neurons. With  $I = 60$  periods of life,

the input layer has dimension  $59 + 5 = 64$ , and the output layer has dimension  $59 + 3 = 62$ . I apply the hyperbolic tangent activation function  $\tanh(\cdot)$  to the hidden layers, which provides smooth, differentiable nonlinearity necessary for backpropagation through the economic equilibrium conditions. Since investment and the predicted synthetic bond price must be nonnegative, I apply the softplus activation function  $\text{softplus}(x) = \log(1 + e^x)$  to the output layer segments corresponding to these quantities, which smoothly enforces the nonnegativity constraint while remaining differentiable everywhere. For equity, all share values must be bounded in the interval  $[0, 1]$  for the market to clear. As such, I apply the sigmoid nonlinear activation function on the segment of the output layer corresponding to equity shares.<sup>3</sup> Lastly, for the equity price I use economic knowledge of the problem to write a custom activation function. Since we know that in the Hayashi (1982) baseline, we should have  $p = QK$  and wedges can drive  $0 < p < QK$ , I evaluate the price of the equity asset as  $p = (1 - \sigma(\cdot))QK \in [0, QK]$  where  $\sigma(\cdot)$  is the sigmoid function taking in the raw FINN output. This specification coaxes the FINN not to find a near-constant solution or one that varies only in the shock process and enforces the economic bounds on the price.

To ensure stable training and keep inputs in the active region of the  $\tanh(\cdot)$  activation function (where gradients are meaningfully nonzero), I normalize all inputs by dividing by  $1 + \bar{\mathcal{X}}$ , where  $\bar{\mathcal{X}}$  is the running mean of the state variable across training iterations. This normalization prevents vanishing gradients and avoids division by zero. The network is trained using the Adam optimizer (Kingma and Ba 2014) with a small learning rate of  $10^{-6}$  to ensure stability in the presence of the complex, nonlinear economic constraints embedded in the loss function. Training proceeds for 5,000 episodes, with minibatches of size 100. Crucially, the simulated training data is redrawn after each epoch, ensuring that the network learns from fresh realizations of the stochastic equilibrium rather than overfitting to a single simulation path. Across all solved models, the loss function evaluates to approximately  $2 \times 10^{-5}$  and the Fisher-Burmeister transform of the Euler equation residual averages approximately 0.35% per household, indicating that the trained networks accurately satisfy the equilibrium conditions.

---

<sup>3</sup>The sigmoid function is given as  $\sigma(x) = \frac{1}{1+e^{-x}}$ . For the sake of full transparency, I adjust this by a ‘temperature’ of 2 so that I actually evaluate  $\sigma(x/2)$ , which helps push some values closer to 0.

Expectations in the Euler equations and forecasts of future state variables are computed using Gauss-Hermite quadrature with 15 nodes. Table 2.2 reports the computational environment and runtime statistics. Training each model to convergence requires approximately 30 minutes on consumer-grade hardware. Unlike machine learning applications that require large pre-existing datasets, the FINN approach generates its own training data through simulation. This means that the memory bottleneck is the FINN size rather than dataset size, allowing the entire training process to fit comfortably within the VRAM limits of consumer GPUs. This demonstrates the computational accessibility of the FINN approach for large-scale stochastic OLG models without requiring specialized high-performance computing infrastructure. I simulate the training data on the CPU (which is faster at this sort of task) and train the FINN (constructing economic quantities and computing loss function, evaluating the parameters step) on the GPU.

Table 2.2: Computational Environment and Performance

| <b>Hardware</b>                     |                                           |
|-------------------------------------|-------------------------------------------|
| Processor                           | 12th Gen Intel i9-12900KF (24) @ 5.100GHz |
| GPU                                 | NVIDIA GeForce RTX 3080 with 8GB of VRAM  |
| RAM                                 | 32 GB                                     |
| <b>Performance</b>                  |                                           |
| Approximate training time per model | 30 minutes                                |
| Training episodes                   | 5,000                                     |
| Minibatch size                      | 100                                       |
| Quadrature nodes                    | 15 (Gauss-Hermite)                        |
| Approximate Euler residual          | 0.35%                                     |
| Approximate Loss function value     | $2 \times 10^{-5}$                        |

We can formally define the FINN parameterized by  $\Theta$  as a function that maps the state variable to the policy functions:

$$V_{\Theta}(\cdot) : \mathcal{X}_t \mapsto \mathcal{Y}_t$$

where  $\mathcal{X}_t = (\{a_{i,t}\}_i, \mathbb{E}_{t-1}[\Lambda_t], p_{t-1}, Q_{t-1}, K_t, Z_t)$  represents the complete state of the economy (all agents' asset holdings including lagged prices and Tobin's Q, firm capital stock, and the aggregate productivity shock) and  $\mathcal{Y}_t = (\{s_{i,t}\}_i, p_t, \mathbb{E}_t[\Lambda_{t+1}], I_t)$  represents the equilibrium policy functions (all

agents' optimal equity share purchases, the equity price, the predicted forward SDF, and firm investment). The neural network  $V_\Theta$  thus approximates the time-homogeneous Markov equilibrium policy functions that characterize the recursive competitive equilibrium defined earlier.

Critically, given the state variable  $\mathcal{X}_t$  and the FINN's predicted policy functions  $\mathcal{Y}_t$ , I can recover all economic quantities needed to evaluate equilibrium conditions using the budget constraints, market clearing conditions, and firm optimality conditions from the formal model. Specifically: the firm's labor demand and wage  $w_t$  are determined by the static first-order condition  $w_t = (1 - \alpha)Y_t/L_t$  given capital stock  $K_t$  and inelastic labor supply; these prices combined with the exogenous efficiency units  $\ell_i$  and predicted equity shares  $s_{i,t}$  determine household incomes and the employer matching contributions  $m_{i,t}$ ; Tobin's  $Q$  is recovered from the firm's Euler equation as  $Q_t = 1 + \Phi_I(I_t, K_t)$ ; and finally the household budget constraints then back out consumption  $c_{i,t}$  for all agents and the firm's budget constraint determines dividends  $d_t$ . This closed-form recovery of all equilibrium objects from  $(\mathcal{X}_t, \mathcal{Y}_t)$  is what allows the loss function to be evaluated purely as a function of the neural network's inputs and outputs, without requiring additional solution steps within each training iteration.

The loss function can be evaluated as

$$\begin{aligned} \mathcal{L}_\Theta(\{\mathcal{X}_t, \mathcal{Y}_t\}_t) = \log \left( 1 + \right. \\ \frac{1}{T(I-1)} \sum_{i,t} \Psi \left( \frac{(u')^{-1} \left( \beta \mathbb{E} \left[ u'(c_{i,t+1}) \cdot \left( 1 + \frac{\partial m_{i,t}}{\partial s_{i,t}} \cdot \frac{p_{t+1} + d_{t+1}}{p_t} \right) \right] \right)}{c_{i,t}} - 1, s_{i,t} \right)^2 \\ + \frac{1}{T} \sum_t \left( \frac{\mathbb{E}[\Lambda_{t+1} (Z_{t+1} F_K(K_{t+1}, L_{t+1}) - \Phi_K(I_{t+1}, K_{t+1}) + (1 - \delta_{t+1})Q_{t+1})]}{Q_t} - 1 \right)^2 \\ + \frac{1}{T} \left( \sum_i (s_{i,t} + m_{i,t}) - 1 \right)^2 + \frac{1}{T} (\mathbb{E}_t^{\text{NN}}[\Lambda_{t+1}] - \mathbb{E}[\Lambda_{t+1}^*])^2 \\ \left. + \frac{1}{T} \sum_t \left( \frac{\sum_i c_{i,t} + I_t + \Phi_t}{Y_t} - 1 \right)^2 \right) \quad (2.66) \end{aligned}$$

where the five terms inside the logarithm correspond to: (i) the mean-square-error across all households

and time periods of the Fisher-Burmeister transformed Euler equation residuals, ensuring that the first-order conditions for household optimization are satisfied with the employer matching wedge  $\partial m_{i,t}/\partial s_{i,t}$  properly embedded; (2) the mean-square-error of the firm's Euler equation for capital investment, ensuring Tobin's Q satisfies the optimality condition; (3) the mean-square-error of the equity market clearing constraint, ensuring all shares are held by households; (4) the mean-square-error between the FINN-predicted synthetic bond price (expected forward SDF) and the Hansen-Jagannathan efficient SDF in expectation; and (5) the mean-square-error of the aggregate resource constraint, ensuring consumption, investment, and adjustment costs sum to output and thereby ruling out pathological solutions where consumption approaches zero. Together, these five components discipline the neural network to find economically plausible solutions that satisfy household and firm optimizations, market clearing, and correctly aggregate heterogeneous household-implied SDFs.

The outer logarithmic transformation  $\log(1 + \cdot)$  serves an important numerical stability purpose during training. In early training iterations when the network parameters are far from equilibrium, the raw loss terms inside the parentheses can be very large, potentially causing gradient explosion and numerical overflow. The logarithmic transformation compresses these large values, preventing the loss from growing unboundedly and ensuring stable gradient-based optimization. Conversely, when the network has converged and the raw loss is very small (near zero), the logarithm satisfies  $\log(1 + x) \approx x$  for small  $x$ , meaning that the transformation becomes approximately linear and does not distort the loss landscape near the optimum. This design ensures stable training throughout the optimization process while preserving sensitivity to small equilibrium violations once the network approaches convergence.

Finally, we have all the pieces in place to define the formal problem solved by the FINN training algorithm:

$$\Theta^* = \arg \min_{\Theta} \mathcal{L}_{\Theta}(\{\mathcal{X}_t, \mathcal{Y}_t\}_t) \quad (2.67)$$

where the set of neural network parameters  $\Theta^*$  will be discovered through the training routine specified in code, usually a modification of stochastic gradient descent such as the Adam optimizer. The trained FINN  $V_{\Theta^*}$  then provides the solutions to the economic model: for any state  $\mathcal{X}_t$ , the equilibrium policy

functions are given by  $\mathcal{Y}_t = V_{\Theta^*}(\mathcal{X}_t)$ , and all other equilibrium objects (consumption, wages, dividends, matching contributions) can be recovered algebraically from the budget constraints and market clearing conditions as described above. By minimizing the loss function—which embeds the economic equilibrium conditions including the employer matching mechanism and firm investment optimality—the FINN training algorithm effectively solves for the recursive Markov equilibrium of the stochastic overlapping generations model with endogenous firm investment and retirement matching.

### 2.5.5 Computation of an Efficient SDF for the Firm’s Euler Equation

Choosing ‘an’ appropriate SDF for use by the firm when discounting future dividends is a difficult task in this incomplete markets model, as the SDF is not unique. Furthermore, constraints faced by households and the wedge resulting from the employer match means that household marginal rates of substitution may not agree.

A numerically stable solution from Hansen and Jagannathan 1991 (HJ) is to pick the most *efficient* SDF laying in the span of the endogenous household intertemporal marginal rates of substitution (MRS) that satisfies the pricing restriction.

At each time  $t$ , let  $\Lambda_{s,t} = \{\Lambda_{s,t,i}\}_i \in \mathbb{R}^{I-1}$  denote the vector of household SDFs (marginal rates of substitution) across  $I$  households for quadrature node indexed by  $s = 1, \dots, S$ :

$$\Lambda_{s,t+1,i+1} = \beta \frac{u'(c_{i+1,t+1}|Z_{t+1} = Z_s)}{u'(c_{i,t})} \cdot \left(1 + \frac{\partial m_{i,t}}{\partial s_{i,t}}\right) \quad (2.68)$$

Let  $R_{s,t}$  denote the gross return on equity at node  $s$ . The quadrature weights are denoted  $\pi = (\pi_1, \dots, \pi_S)$  with  $\sum_s \pi_s = 1$ . Define the *pricing vector* as

$$b_t = \mathbb{E}_\pi[\Lambda_{s,t} R_{s,t}] \in \mathbb{R}^{I-1} \quad (2.69)$$

Define the second-moment matrix of household SDFs:

$$\Sigma_t = \sum_{s=1}^S \pi_s \Lambda_{s,t} \Lambda'_{s,t} \in \mathbb{R}^{(I-1) \times (I-1)} \quad (2.70)$$

$$\bar{\Sigma}_t = \Sigma_t + \gamma_t \mathbb{I}, \quad \gamma_t = \frac{10^{-4} \text{Tr}(\Sigma_t)}{I} \quad (2.71)$$

To ensure numerical stability (since  $\Sigma_t$  is low rank when the number of quadrature nodes is less than the number of households), a small ridge  $\gamma_t$  in the diagonals ( $\mathbb{I}$  being the identity matrix) is added to regularize before inverting the second-moment matrix.

The HJ problem is

$$\min_{\omega_t} \left\{ \frac{1}{2} \omega'_t \bar{\Sigma}_t \omega_t \right\} \quad \text{subject to} \quad \mathbb{E}_\pi [\omega'_t \Lambda_{s,t} R_{s,t}] = 1. \quad (2.72)$$

The solution can be given in closed form:

$$\omega_t^* = \frac{(\bar{\Sigma}_t)^{-1} b_t}{b'_t (\bar{\Sigma}_t)^{-1} b_t} \in \mathbb{R}^{I-1} \quad (2.73)$$

Finally, the implied *firm SDF* at each state  $(s, t)$  is

$$\Lambda_{s,t}^* = (\omega_t^*)' \Lambda_{s,t} \in \mathbb{R} \quad (2.74)$$

By construction, the projection enforces the asset pricing restriction within the span of household SDFs:

$$\mathbb{E}_\pi [\Lambda_t^* R_t] = 1, \quad \Lambda_{s,t}^* \in \text{span} \{ \Lambda_{s,t} \} \quad (2.75)$$

meaning that the SDF used by the firm is a linear combination of endogenous household MRS and therefore accurately reflects household pricing, while the efficiency of the SDF (minimizing the second moment) means that it remains computationally stable and is robust to abrupt changes to the MRS of individual households (which are common in early training).

### 2.5.6 Results

The numerical results confirm the analytical predictions from the two-period model and reveal substantial quantitative effects of employer matching on asset prices, real outcomes, and welfare. I simulate the economy under three matching regimes: no matching ( $\psi = 0$ ), 50-cent-on-the-dollar matching up to 6% of income ( $\psi = 0.5, \phi = 0.06$ ), and dollar-for-dollar matching up to 6% of income ( $\psi = 1, \phi = 0.06$ ). These parameterizations correspond to common employer matching policies observed in practice. For each regime, I compute ergodic distributions by simulating trajectories of 100,000 periods post-training and report unconditional means.

Table 2.3 presents the asset pricing outcomes. The central mechanism of the paper operates precisely as predicted: employer matching raises the SDF and lowers equilibrium equity returns. Moving from no matching to dollar-for-dollar matching increases the expected SDF by 68 basis points (from 0.9168 to 0.9237), reflecting that households become more patient when facing stronger incentives to save in order to capture employer contributions. This increase in the SDF directly translates into lower equilibrium returns. The average equity return falls by 79 basis points (from 9.08% to 8.29%) when introducing dollar-for-dollar matching, with the 50-cent match producing an intermediate decline of 47 basis points. The mechanism here is transparent: households are willing to accept lower market returns because their effective returns—inclusive of the employer match subsidy—remain attractive. From the household's perspective, even though market returns have fallen, the match subsidy more than compensates, making equity purchases worthwhile. Importantly, this return decline is not a mechanical consequence of reduced dividends, but rather reflects the endogenous adjustment of the pricing kernel.

The lower cost of capital induced by employer matching stimulates firm investment and raises the aggregate capital stock, as shown in Table 2.4. With dollar-for-dollar matching, the capital stock increases by 6.1% relative to the no-matching baseline (from 19.181 to 20.345). This increased capital accumulation reflects the firm's optimal response to facing a higher SDF: when future dividends are discounted less heavily, the firm finds it profitable to increase investment in physical capital. The 50-cent match

Table 2.3: Asset Pricing Outcomes Across Economies ( $\psi$  values)

|                          | $\psi = 0$ | $\psi = 0.5$ | $\psi = 1$ |
|--------------------------|------------|--------------|------------|
| Average return (%)       | 9.08       | 8.61         | 8.28       |
|                          | –          | -47 bps      | -79 bps    |
| Volatility (%)           | 0.54       | 0.57         | 0.70       |
|                          | –          | +3 bps       | +16 bps    |
| $\mathbb{E}[\text{SDF}]$ | 0.9168     | 0.9208       | 0.9237     |
|                          | –          | +39 bps      | +68 bps    |

produces an intermediate capital increase of 3.5%, demonstrating that the effect scales smoothly with matching generosity. The increase in capital stock raises output through the Cobb-Douglas production function: dollar-for-dollar matching increases output by 1.7% (from 12.167 to 12.379), while the 50-cent match raises output by 1.0%. Since labor is supplied inelastically at  $L = 1$  and capital and labor are complements in production, the higher capital stock raises the marginal product of labor proportionally, increasing equilibrium wages by the same 1.7% for dollar-for-dollar matching (from 8.517 to 8.665) and 1.0% for the 50-cent match. These wage gains represent a general equilibrium spillover from the matching policy: all workers benefit from higher wages due to capital levels rising, not just those who directly receive employer matches. This finding suggests that employer matching may have progressive distributional consequences beyond the direct transfer to participants.

Table 2.4: Real Outcomes Across Economies ( $\psi$  values)

|         | $\psi = 0$ | $\psi = 0.5$ | $\psi = 1$ |
|---------|------------|--------------|------------|
| Output  | 12.167     | 12.291       | 12.379     |
|         | –          | +1.0%        | +1.7%      |
| Capital | 19.181     | 19.856       | 20.345     |
|         | –          | +3.5%        | +6.1%      |
| Wage    | 8.517      | 8.604        | 8.665      |
|         | –          | +1.0%        | +1.7%      |

Table 2.5 reports consumption-equivalent welfare gains for a newborn household entering the economy. Dollar-for-dollar matching generates a 2.6% consumption-equivalent welfare gain relative to the no-matching baseline, while 50-cent matching produces a 0.3% gain. These welfare gains reflect the interplay of several economic forces. Employer matching induces households to tilt their savings profiles

earlier in the life cycle, purchasing equity shares at younger ages when employer contributions subsidize these purchases. This earlier tilting of the savings profile is directly consistent with the increase in the SDF documented in Table 2.3: matching effectively simulates increased patience, making households behave as if they are more forward-looking and therefore finding it optimal to save earlier in life to capture subsidized returns. Since the total number of equity shares is fixed at one, this intertemporal reallocation of share ownership means that households with matching hold a larger fraction of the firm earlier in life, when their marginal utility of consumption is higher and their discount factor weights consumption more heavily. This front-loading of equity ownership enables households to finance steeper consumption profiles that deliver more utility in earlier, less-discounted periods of life.

Table 2.5: Consumption-Equivalent Welfare Across Economies ( $\psi$  values)

|                           | $\psi = 0$ | $\psi = 0.5$ | $\psi = 1$ |
|---------------------------|------------|--------------|------------|
| Expected Lifetime Utility | -30.271    | -30.229      | -29.926    |
| CEW                       | –          | +0.3%        | +2.6%      |

The substantial positive welfare gains reported here indicate that the intertemporal consumption smoothing benefits and wage spillovers more than compensate for the return reduction, particularly for newborn households who have long horizons over which to benefit from the tilted savings profile and higher wages. The convexity of the welfare gains with respect to matching intensity (0.3% for  $\psi = 0.5$  but 2.6% for  $\psi = 1$ ) suggests that more generous matching policies generate disproportionately large welfare improvements, potentially due to nonlinearities in how matching incentives interact with life-cycle savings behavior.

Figures 2.2, 2.3, and 2.4 display the expected life-cycle profiles of the matching wedge, savings, and consumption across matching regimes, revealing how households adjust their intertemporal decisions in response to matching incentives. Figure 2.2 shows the matching wedge  $\partial m_i / \partial s_i$  that appears in the household Euler equation, representing the marginal subsidy to equity purchases. The wedge is identically zero throughout life for the no-matching economy, confirming that households face no distortion to their savings decisions. With matching, the wedge is largest at the beginning of life and declines monotonically




Figure 2.2: Expected Life-Cycle Matching Wedge  $\partial m_i / \partial s_i$  Across Economies ( $\psi$  values)




Figure 2.3: Expected Life-Cycle Savings Profile Across Economies ( $\psi$  values)

over the life cycle, reaching near-zero by around age 40 (period 20) as households accumulate sufficient assets that their marginal contributions are above the matching cap. This monotonically decreasing pattern reflects that young households with low asset stocks have marginal contributions below the cap and therefore receive the full marginal match subsidy on additional equity purchases, while older households with accumulated wealth find their marginal contributions constrained by the cap, causing the matching wedge to disappear from their Euler equations. Importantly, the wedge is uniformly larger for dollar-for-dollar matching than for the 50-cent match at every age, showing that more generous matching policies create stronger incentives to shift savings earlier in the life cycle. This front-loaded, age-varying subsidy structure directly explains the steepening of early-life savings documented in Figure 2.3.

Figure 2.3 shows that employer matching induces households to accumulate equity earlier in the life cycle and sustain higher savings throughout working years, precisely when the matching wedge is active.




Figure 2.4: Expected Life-Cycle Consumption Profile Across Economies ( $\psi$  values)

In the no-matching economy, the savings profile exhibits the standard hump shape: households gradually build up retirement wealth during their working years, peak around age 50-55, and then decumulate in retirement. With matching, this profile shifts upward and becomes steeper, particularly during the early working years when the matching wedge in Figure 2.2 is largest. The effect is most pronounced for dollar-for-dollar matching, where households accumulate substantially more wealth early in life to maximize the capture of employer contributions while the wedge remains active. This pattern confirms the mechanism highlighted in the introduction: households tolerate lower equilibrium returns because they are effectively earning higher returns inclusive of the match subsidy shown in Figure 2.2, making it optimal to save more despite the general equilibrium price adjustment. Figure 2.4 shows the corresponding consumption profiles. Together, these life-cycle profiles illustrate that employer matching fundamentally reshapes the timing of savings and consumption decisions, with households responding to the front-loaded matching wedge by accelerating savings during early working years and enjoying higher lifetime consumption as a result of both direct transfers and general equilibrium wage spillovers.

## 2.6 Conclusion

This paper studies the general equilibrium effects of employer retirement matching on firm investment, asset prices, and macroeconomic outcomes. The central premise is that employer matching contributions subsidize workers' equity purchases, creating a wedge in the household Euler equation that increases the SDF. As households become effectively more patient—eager to capture employer contributions that would otherwise be left on the table—they tolerate lower equilibrium returns while still earning attractive effective returns inclusive of the match subsidy. This increase in the SDF reduces the firm's cost of capital, stimulating investment in physical capital and raising aggregate output and wages. The key question is whether this mechanism generates meaningful macroeconomic effects or merely redistributes resources from shareholders to workers.

I establish the core mechanism analytically in a two-period deterministic model without aggregate risk. The closed-form results demonstrate that employer matching unambiguously increases the SDF, reduces equilibrium equity returns, and raises the capital stock. These analytical predictions hold under both dividend-financed matching (the baseline specification) and a robustness check with worker-financed matching, confirming that the results do not arise mechanically from reduced dividend payments but rather reflect the endogenous adjustment of the pricing kernel through household savings behavior. The two-period model provides transparent intuition for the economic forces at work and shows that the qualitative effects are robust features of the model structure.

To quantify the macroeconomic importance of employer matching, I solve a sixty-period stochastic overlapping generations model paired with a neoclassical model of firm investment featuring convex adjustment costs and aggregate productivity risk. The computational challenge lies in solving for equilibrium policy functions in a high-dimensional state space with heterogeneous households whose marginal rates of substitution differ due to binding short-sale constraints and the matching wedge. I employ finance-informed neural networks (FINNs) to approximate the global policy functions, embedding household Euler equations, firm optimality conditions, and market clearing constraints directly into the loss func-

tion. To aggregate heterogeneous household SDFs into a single endogenous discount rate for the firm's investment problem, I construct the Hansen-Jagannathan efficient SDF that lies in the span of household marginal rates of substitution while satisfying the asset pricing restriction. This approach ensures computational stability and accurately reflects household pricing despite heterogeneity in savings behavior.

The numerical results confirm and quantify the analytical predictions. Introducing dollar-for-dollar employer matching at realistic parameter values reduces equilibrium equity returns by 79 basis points, increases the aggregate capital stock by 6.1%, and raises output and wages by 1.7%. The effects scale smoothly with matching generosity: 50-cent-on-the-dollar matching produces intermediate effects of 47 basis points lower returns and 3.5% higher capital. Life-cycle profile analysis reveals that the matching wedge  $\partial m / \partial s$  is largest early in life and declines monotonically as households accumulate assets and their marginal contributions rise above the matching cap, causing the wedge to vanish from the Euler equation. This front-loaded subsidy structure induces households to accelerate equity purchases during early working years, tilting savings profiles earlier and enabling steeper consumption profiles that deliver more utility in less-discounted periods. Newborn households entering an economy with dollar-for-dollar matching experience consumption-equivalent welfare gains of 2.6% relative to a no-matching baseline, reflecting the combined benefits of intertemporal consumption smoothing and general equilibrium wage spillovers that more than offset lower equilibrium returns.

These findings demonstrate that employer matching has first-order macroeconomic effects that extend well beyond simple redistribution between shareholders and workers. By subsidizing household equity demand, matching increases the SDF and lowers the cost of capital, fundamentally reshaping the investment opportunity set available to firms. The resulting capital deepening raises wages for all workers—not just those receiving matches—illustrating a general equilibrium spillover channel through which retirement policy affects macroeconomic outcomes. More broadly, this paper contributes to the literature on the real effects of financial frictions and household portfolio choice by showing that seemingly micro-level retirement policies can have substantial aggregate implications when analyzed in general equi-

librium. Future research could usefully extend this framework to incorporate heterogeneity in matching generosity across firms and income groups, endogenize the firm's choice of matching policy through bargaining or tax incentives, and study the transition dynamics as the economy moves between matching regimes.

# Chapter 3: Deep Learning the Term Structure for Derivatives Pricing

## 3.1 Introduction

The Heath-Jarrow-Morton (HJM) framework (Heath, Jarrow, and Morton 1992) is the most general arbitrage-free approach to pricing interest rate derivatives, specifying stochastic evolution for the entire forward curve simultaneously and nesting earlier short-rate models as special cases. Under risk-neutral valuation, a derivative with contract features  $\Xi$  (including expiry  $T$ ) and payoff at expiry  $h(T, f; \Xi)$  depending on the forward curve  $f(t, T)$  has time- $t$  price given by the stochastic expectation:

$$V(t) = \mathbb{E} \left[ \exp \left( - \int_t^T r(s) ds \right) h(T, f; \Xi) \right] \quad (3.1)$$

where  $r(t) = f(t, t)$  is the short rate. Path dependence enters through two channels: the discount factor requires integrating the short rate  $r(s)$  along the entire stochastic path of the forward curve from  $t$  to  $T$ , and the payoff  $h(T, f; \Xi)$  depends on the forward curve realized at expiry. This path dependence prevents closed-form analytical solutions for most derivatives, forcing practitioners to rely on numerical methods. Since the forward curve is infinite-dimensional, achieving tight estimates of derivative prices entails evaluating this expectation over a fine partition of the forward curve. Monte Carlo simulation of thousands or millions of stochastic forward curve paths becomes increasingly computationally expensive as this grid grows finer, forcing practitioners to choose between accuracy and computational feasibility.

Suppose you simulate a contract and obtain a price with reasonable accuracy. If the counterparty changes the terms—even marginally adjusting strike, maturity, or other contract terms—you must resimulate entirely from scratch. Worse, computing Greeks requires perturbing contract parameters or underlying rates and drawing fresh paths for each sensitivity, making real-time risk management computationally expensive.

This paper introduces *Finance-Informed Neural Networks* (FINNs) to solve this computational bottleneck. FINNs circumvent Monte Carlo entirely by exploiting the Feynman-Kac theorem, which establishes that the stochastic expectation above satisfies a deterministic PDE that can be solved directly via deep learning. The approach combines two key insights: first, the Feynman-Kac transformation eliminates Monte Carlo simulation by replacing the stochastic pricing problem with a PDE characterization; second, neural networks trained via automatic differentiation solve high-dimensional PDEs efficiently, avoiding the curse of dimensionality that plagues traditional finite difference methods. Once trained, FINNs price derivatives in microseconds regardless of state space dimension, delivering speedups of many orders of magnitude over Monte Carlo. Critically, the major Greeks—theta and curve deltas—come for free, since they appear directly in the PDE being minimized during training. Other Greeks require only negligible additional computation via automatic differentiation. A further practical advantage is the approach’s flexibility: regardless of the specific derivative contract—caplets, swaptions, callable bonds, or exotic path-dependent structures—the core PDE governing prices remains identical. Pricing different instruments requires only adjusting the contract features in the state variable and modifying the boundary condition to reflect the appropriate payoff function, without altering the fundamental PDE structure.

The methodology builds on *physics-informed neural networks* (PINNs) pioneered by Raissi, Perdikaris, and Karniadakis (2019), which embed governing differential equations directly into the loss function. Neural networks are universal function approximators whose evaluation cost scales gracefully with dimension. Modern deep learning frameworks compute exact derivatives via automatic differentiation, enabling direct evaluation of PDE residuals at any point in the state space. The neural network param-

eterizes the pricing functional, and training minimizes PDE violations alongside boundary condition penalties—critically, enforced without any forward simulation by evaluating terminal payoffs directly on the cross-section of historical forward curves rather than simulating curves to maturity.

I demonstrate this methodology by pricing interest rate caplets—call options on future LIBOR rates and fundamental building blocks for caps and swaptions. Using daily U.S. Treasury forward curve data from Gürkaynak, Sack, and Wright (2007), I estimate volatility via PCA, derive the caplet pricing PDE, and train neural networks across eight discretization levels ranging from  $K = 10$  to  $K = 150$  tenor points. The training procedure exploits the analytical zero-strike caplet solution as an additional anchor, disciplining the pricing functional with exact solutions where available. All results use local volatility (scaling with the square root of forward rate levels) to capture the empirical phenomenon that interest rate volatility scales with rate levels. Validated against a test set of 1,000 randomly sampled contracts, the trained FINNs achieve pricing accuracy within 0.04¢ to 0.07¢ compared to Monte Carlo benchmarks (per dollar of contract value), while requiring only consumer-grade hardware (8GB GPU) and evaluating in a few microseconds once trained. The computational advantage is dramatic: as I grow the state space (discretization of the forward curve) from 10 to 150 nodes, FINNs price caplets 300,000 to 4,500,000 times faster than Monte Carlo simulation.

## 3.2 Literature Review

This paper sits at the intersection of term structure modeling, computational PDE methods, and machine learning for derivatives pricing.

Early short-rate models (Vasicek 1977; Cox, Ingersoll, Ross, et al. 1985; Hull and White 1990) specify dynamics for the instantaneous rate and derive the entire term structure from this single state variable, offering analytical tractability at the cost of restricting all interest rate movements to a single factor. For interest rate derivatives, practitioners frequently use Black (1976) formula for caps and floors, which applies the Black-Scholes framework directly to forward rates. However, the Black model prices each caplet independently without specifying how forward rates co-evolve, potentially admitting arbitrage when pricing portfolios of caplets at different maturities.

Heath, Jarrow, and Morton (1992) revolutionized the field by modeling the entire forward curve simultaneously, yielding an arbitrage-free framework that nests earlier short-rate models as special cases. The cost of this generality is computational: the forward curve is infinite-dimensional, requiring discretization for numerical implementation. Subsequent work has taken two approaches to manage this complexity. One strand imposes structure to achieve dimensionality reduction, either through Markovian restrictions (Cheyette 2001) or through discretely-compounded forward rate specifications like the LIBOR market model (Brace, Gatarek, and Musiela 1997), which prices caps and swaptions directly in terms of observable market rates. My approach takes a different path, avoiding dimensionality reduction entirely—the neural network accommodates the full discretized forward curve without requiring Markovian structure or parametric restrictions. The Musiela parameterization (Musiela and Rutkowski 2005) reformulates HJM dynamics in terms of time-to-maturity rather than calendar time, a change of variables essential for PDE-based pricing methods.

PDEs have been understood to be at the heart of option pricing at least since the seminal work of Black and Scholes (1973) and Merton et al. (1971). Recent work applies neural networks to solve pricing PDEs in continuous-time models. Gopalakrishna (2021) and A. Zhang (2022) demonstrate neural network

solutions for PDEs arising in macro-finance models, while a parallel literature (Chen et al. [2023](#); Yuana [2024](#); Cao, Chen, and Hull [2020](#); Cao, Chen, Hull, and Poulos [2021](#); Cao, Chen, Farghadani, et al. [2023](#)) focuses on learning optimal hedging strategies and/or derivatives prices.

A distinct methodological strand uses neural networks to *solve* PDEs without discretizing the state space. Raissi, Perdikaris, and Karniadakis ([2019](#)) pioneered physics-informed neural networks (PINNs), which embed the governing PDE directly in the loss function via automatic differentiation, eliminating the need for spatial grids. This mesh-free approach scales naturally to high dimensions. Han, Jentzen, and E ([2018](#)) develop the deep BSDE method for backward stochastic differential equations, demonstrating solutions to parabolic PDEs in hundreds of dimensions, while Sirignano and Spiliopoulos ([2018](#)) propose the Deep Galerkin Method as an alternative variational approach. For financial applications specifically, Beck et al. ([2021](#)) extend these methods to Kolmogorov PDEs, solving backward in time from terminal conditions.

### 3.3 Heath-Jarrow-Morton Model Refresher

What follows is a non-rigorous refresher of the Heath-Jarrow-Morton (Heath, Jarrow, and Morton 1992) paper.

Consider a continuous trading interval  $[0, \bar{T}]$  for a fixed  $\bar{T} > 0$  and probability space  $(\Omega, \mathcal{F}, Q)$  where  $\Omega$  is the state space,  $\mathcal{F}$  is the  $\sigma$ -algebra representing measurable events, and  $Q$  is a probability measure. Augmented, right-continuous complete filtration  $\{F_t : t \in [0, \bar{T}]\}$  generated by  $N \geq 1$  independent Brownian motions  $\{W_n(t) : t \in [0, \bar{T}]\}_{n=1}^N$  initialized at zero. Let  $\mathbb{E}[\cdot]$  denote expectations with respect to the probability measure  $Q$ .

There exists a continuum of default-free pure discount bonds trading with differing maturities  $T \in [0, \bar{T}]$ . The price at time  $t$  of a bond maturing at  $T$  for all  $T \in [0, \bar{T}]$  and  $t \in [0, T]$  is denoted  $P(t, T)$ . Face values are normalized to 1:  $P(T, T) = 1$  for all  $T \in [0, \bar{T}]$ . Additionally,  $P(t, T) > 0$  and  $\partial \log P(t, T) / \partial T$  exists for all  $T \in [0, \bar{T}]$  and  $t \in [0, T]$ .<sup>1</sup>

Define the instantaneous forward rate at time  $t$  for date  $T > t$  as:

$$f(t, T) = -\frac{\partial \log P(t, T)}{\partial T} \text{ for all } T \in [0, \bar{T}], t \in [0, T]. \quad (3.2)$$

Bond prices can be expressed in terms of forward rates:

$$P(t, T) = \exp \left( - \int_t^T f(t, s) ds \right) \text{ for all } T \in [0, \bar{T}], t \in [0, T]. \quad (3.3)$$

The instantaneous forward rate is termed the *spot rate* and is given by

$$r(t) = f(t, t) \text{ for all } t \in [0, \bar{T}]. \quad (3.4)$$

---

<sup>1</sup>Normalize payoffs, no arbitrage, and forward rates are well-defined, respectively.

HJM starts with forward rate dynamics given as:

$$df(t, T) = \mu(t, T)dt + \sigma(t, T) \cdot dW(t). \quad (3.5)$$

Above,  $\sigma(t, T) = (\sigma_n(t, T))_{n=1}^N$  and  $W(t) = (W_n(t))_{n=1}^N$  are vectors in  $\mathbb{R}^n$  representing the  $n \geq 1$  volatility processes. Both  $\mu$  and  $\sigma$  are assumed to be measurable, adapted, and integrable over  $[0, T]$  almost everywhere with respect to  $Q$ .

$$\begin{aligned} \int_0^T |\mu(t, T)| dt &< +\infty \\ \int_0^T \sigma_n^2(t, T) dt &< +\infty \text{ for } i = n, \dots, N. \end{aligned}$$

Given a deterministic initial forward curve  $f(0, T)$ , the dynamics above uniquely determine the stochastic fluctuation of the entire forward curve according to

$$f(t, T) = f(0, T) + \int_0^t \mu(s, T)ds + \sum_{n=1}^N \int_0^t \sigma_n(s, T)dW_n(s) \quad (3.6)$$

for all  $0 \leq t \leq T$ .

Applying Itô's Lemma to equation (3.3) and (3.5) and applying the no-arbitrage restriction yields the familiar HJM result:<sup>2</sup>

$$\mu(t, T) = \sigma(t, T) \cdot \int_t^T \sigma(t, s)ds \quad (3.7)$$

This result implies that the dynamics of the entire forward curve are parameterized only by choice of  $\sigma_i$  functions. The next section will discuss historical methods for doing so, as well as common methods for using this model for contingent claims valuation. The following section will discuss an improved contingent claims valuation approach, blending deep learning with the Feynman-Kac theorem.

---

<sup>2</sup>For a detailed derivation, see the original paper.

### 3.3.1 Musiela Parameterization

In computational applications, a change of variables termed the ‘Musiela Parameterization’ is undertaken: from  $(t, T)$  to  $(t, \tau)$  where  $\tau = T - t$  is termed the *tenor* and represents time-to-maturity (Musiela and Rutkowski 2005). The model dynamics then become:

$$df(t, \tau) = \mu(t, \tau)dt + \sigma(t, \tau) \cdot dW_t \quad (3.8)$$

$$d\tau = -dt \quad (3.9)$$

and the state space is  $t \in [0, \bar{T}]$ ,  $\tau \in [0, \bar{T} - t]$ .

The no-arbitrage drift process is also amended:

$$\mu(t, \tau) = \frac{\partial}{\partial \tau} f(t, \tau) + \sigma(t, \tau) \cdot \int_0^\tau \sigma(t, \tau s) ds \quad (3.10)$$

The Musiela form will be used in what follows for the rest of the paper.

### 3.4 Standard Procedure for Computing the HJM Model

The standard approach to implementing the HJM model involves three computationally straightforward steps: obtaining forward curve data, estimating the volatility structure via principal components analysis, and computing the no-arbitrage drift via numerical integration. All three steps are relatively inexpensive and can be performed once as preprocessing. I follow this standard procedure in my implementation, using widely available data and conventional estimation techniques.

The computational expense in HJM modeling does not arise from these preprocessing steps. Rather, it emerges when pricing path-dependent interest rate derivatives, which traditionally requires Monte Carlo simulation of thousands or millions of forward curve paths. This Monte Carlo bottleneck is the problem I address in subsequent sections by first casting the *stochastic* simulation exercise into a *deterministic* PDE and second by using FINNs to solve the pricing PDE directly.

#### 3.4.1 Data and Forward Curve Construction

This paper uses daily instantaneous forward rate data constructed by Gürkaynak, Sack, and Wright [2007](#) from U.S. Treasury securities. The dataset provides fitted Svensson parameters  $(\beta_0, \beta_1, \beta_2, \beta_3, \tau_1, \tau_2)$  estimated daily from all traded Treasury securities using the parametric yield curve methodology of Svensson [1994](#). I utilize data from January 1, 2001 onward, chosen to focus on the modern interest rate environment while maintaining a sufficiently long time series for estimating volatility dynamics.

The Svensson model specifies the instantaneous forward rate as a function of tenor:

$$f(\tau) = \beta_0 + \left( \beta_1 + \beta_2 \frac{\tau}{\tau_1} \right) \exp \left( -\frac{\tau}{\tau_1} \right) + \beta_3 \frac{\tau}{\tau_2} \exp \left( -\frac{\tau}{\tau_2} \right) \quad (3.11)$$

Raw Svensson parameter estimates occasionally exhibit extreme values due to market stress or thin trading in certain maturities. To ensure numerical stability, I filter observations using quantile-based outlier detection: for each parameter, I retain only observations between the 5th and 95th percentiles of its empirical

distribution. This removes approximately 10% of observations while preserving the full range of typical market conditions. For implementation, I evaluate the Svensson formula at a grid of equally spaced intervals over the interval  $[0, 5]$  years—appropriate for pricing short-maturity caplets. A final filtering step removes any dates where forward rates are negative or numerically close to zero (below  $\epsilon = 0.005 = 0.5\%$ ) at any tenor.

### 3.4.2 Volatility Estimation

Practitioners discretize the tenor dimension into a finite grid  $\{\tau_k\}_{k=1}^K$  and estimate the covariance matrix of forward rate changes across these tenors. Because this covariance is computed by pooling data across time, the time dimension  $t$  is integrated out, yielding time-invariant volatility estimates. I use the Federal Reserve’s published forward rates at annual tenor intervals for  $\tau = 1, \dots, 30$  years for this volatility estimation.

I compute the sample covariance matrix of daily forward rate changes, scaled to annual units by multiplying by 252 trading days:

$$\text{Cov}[\Delta f] = 252 \cdot \mathbb{E}[(\Delta f)(\Delta f)'] \quad (3.12)$$

where  $\Delta f$  denotes the vector of daily changes in forward rates across the 30 tenor points.

Principal components analysis (PCA) is then applied to this covariance matrix to extract the dominant factors driving term structure movements. Empirically, the first three principal components typically capture a large majority of the variation in forward rate changes, corresponding to the well-known level, slope, and curvature factors. Applying eigenvalue decomposition to the covariance matrix, I extract the three eigenvectors corresponding to the three largest eigenvalues. To obtain volatility magnitudes, I scale each eigenvector by the square root of its corresponding eigenvalue, yielding adjusted volatility vectors:

$$\tilde{\sigma}_n(\tau_k) = \sqrt{\lambda_n} \cdot v_n(\tau_k) \quad (3.13)$$

where  $\lambda_n$  is the  $n$ -th largest eigenvalue and  $v_n$  is the corresponding eigenvector, evaluated at the discrete

grid points  $\tau_k = k$  for  $k = 1, \dots, 30$ .

To obtain volatility functions  $\sigma_n(\tau)$  defined over the continuum of tenors, I fit smooth Chebyshev polynomials through these discrete principal component loadings. Specifically, for each factor  $n = 1, 2, 3$ , I fit a Chebyshev polynomial of degree 3:

$$\sigma_n(\tau) = \sum_{j=0}^3 c_{n,j} T_j \left( \frac{2\tau}{\tau_{\max}} - 1 \right) \quad (3.14)$$

where  $T_j$  is the  $j$ -th Chebyshev polynomial of the first kind and  $c_{n,j}$  are fitted coefficients.

### 3.4.3 Computing the No-Arbitrage Drift

A key advantage of the Svensson parameterization is that it provides an analytical expression for the derivative of the forward curve with respect to tenor:

$$\frac{df}{d\tau} = \left( -\frac{\beta_1}{\tau_1} + \frac{\beta_2}{\tau_1} \left( 1 - \frac{\tau}{\tau_1} \right) \right) \exp\left( -\frac{\tau}{\tau_1} \right) + \frac{\beta_3}{\tau_2} \left( 1 - \frac{\tau}{\tau_2} \right) \exp\left( -\frac{\tau}{\tau_2} \right) \quad (3.15)$$

With the volatility structure  $\sigma_n(\tau)$  and forward curve derivative  $\partial f / \partial \tau$  specified, the no-arbitrage drift can be computed from the Musiela condition (equation 101):

$$\mu(t, \tau) = \frac{\partial}{\partial \tau} f(t, \tau) + \sigma(t, \tau) \cdot \int_0^\tau \sigma(t, s) ds \quad (3.16)$$

The second term requires evaluating the integral  $\int_0^\tau \sigma(t, s) ds$  for each volatility factor. Since the volatility functions are represented as Chebyshev polynomials, this integral is computed numerically using standard quadrature methods such as the trapezoidal rule or Simpson's rule over the discretized tenor grid  $\{\tau_k\}_{k=1}^K$ .

This numerical integration step is computationally inexpensive and needs to be performed only once per forward curve evaluation. Crucially, the drift computation does not require stochastic simulation—it is a deterministic calculation given the current state of the forward curve  $f(t, \tau)$  and the pre-estimated

volatility structure  $\sigma_n(\tau)$ .

These fitted volatility functions  $\sigma_n(\tau)$  and the associated drift will be held fixed during neural network training, entering the loss function through the drift term and the second-order derivative terms in the PDE.

### 3.4.4 Local Volatility Specification

The volatility estimation procedure described above yields time-invariant volatility functions  $\sigma_n(\tau)$  that depend only on tenor. However, I adopt a local volatility specification that allows volatility to depend on the current level of forward rates. Local volatility models are widely used in interest rate derivatives markets to capture the empirical phenomenon that interest rate volatility tends to scale with the level of rates—a feature not present in constant-volatility specifications. All results presented in this paper use this local volatility structure.

The Finance-Informed Neural Network approach handles local volatility with minimal additional complexity. State-dependent volatility enters the model simply through automatic differentiation when computing the PDE residual, requiring no changes to the neural network architecture or training procedure. This flexibility stands in stark contrast to Monte Carlo methods, where local volatility substantially increases computational burden: each simulated path must now track both the forward curve evolution and the state-dependent volatility at each time step, with the volatility function evaluated thousands of times per simulation.

For the FINN implementation, the only modification required is in the data preprocessing step. I compute the covariance matrix using proportional changes rather than absolute changes:

$$\text{Cov} \left[ \frac{\Delta f}{\sqrt{f}} \right] = 252 \cdot \mathbb{E} \left[ \left( \frac{\Delta f}{\sqrt{f}} \right) \left( \frac{\Delta f}{\sqrt{f}} \right)' \right] \quad (3.17)$$

where the division by  $\sqrt{f}$  scales each forward rate change by the square root of the current forward rate level. After computing the scaled covariance matrix, I apply the same PCA decomposition and Cheby-

shev polynomial fitting described in the previous subsection. The resulting volatility functions  $\sigma_n(\tau)$  now represent proportional volatilities.

During neural network training, these proportional volatilities are multiplied by  $\min\{\sqrt{f(t, \tau)}, M\}$  where  $M = 0.4 = 40\%$  is the parameter suggested in the seminal continuous-time textbook from Shreve (2004). The cap  $M$  prevents numerical instability when forward rates become extremely large—limiting the volatility scaling to reasonable levels even in high-rate environments. This choice ensures that the volatility scaling remains well-behaved across the full range of observed forward rate environments. With local volatility, the HJM dynamics under the Musiela parameterization become:

$$df(t, \tau) = \mu(t, \tau)dt + \sigma(t, \tau, f) \cdot dW_t \quad (3.18)$$

$$\mu(t, \tau) = \frac{\partial}{\partial \tau} f(t, \tau) + \sigma(t, \tau, f) \cdot \int_0^\tau \sigma(t, s, f) ds \quad (3.19)$$

where the state-dependent volatility function is defined as:

$$\sigma(t, \tau, f) = \tilde{\sigma}(\tau) \cdot \min\{\sqrt{f(t, \tau)}, M\} \quad (3.20)$$

with  $\tilde{\sigma}(\tau)$  denoting the PCA-derived proportional volatility functions computed during preprocessing. Importantly, the PCA-derived volatility values  $\tilde{\sigma}_n(\tau)$  are computed only once during the preprocessing stage and stored as fixed parameters. During neural network training, multiplying these stored values by  $\min\{\sqrt{f(t, \tau)}, M\}$  to produce the state-dependent volatility is computationally negligible.

With all preprocessing complete—forward curve data obtained, volatility structure estimated, and drift computation specified—the standard HJM implementation would proceed to Monte Carlo simulation for pricing path-dependent derivatives. My method avoids this computational bottleneck entirely. To explain how, I require a brief mathematical digression to the Feynman-Kac theorem, which transforms the stochastic pricing problem into a deterministic partial differential equation. Once this theoretical foundation is established, I will present the specific application to caplet pricing and describe the neural network algorithm that solves the PDE.

### 3.5 Feynman-Kac Formula in the HJM Framework

The Feynman-Kac theorem provides the mathematical foundation for my approach to pricing interest rate derivatives. This classical result from stochastic analysis establishes an equivalence between stochastic expectations and deterministic partial differential equations. The key insight is that instead of computing the expected discounted payoff through Monte Carlo simulation—which requires generating thousands or millions of random paths—I can solve a PDE that characterizes the same pricing functional. This section presents the multidimensional Feynman-Kac theorem and shows how to apply it to the HJM framework.

Recall the fundamental pricing formula under risk-neutral valuation. For a derivative security with contract features  $\Xi$  and payoff function  $h(t, f; \Xi)$  depending on the forward curve, the time- $t$  price is given by:

$$V(t, \tau, f; \Xi) = \mathbb{E} \left[ \exp \left( - \int_t^{t+\tau} r(s) ds \right) h(t + \tau, f; \Xi) \mid \mathcal{F}_t \right] \quad (3.21)$$

where  $\tau$  is the time-to-maturity,  $r(s) = f(s, 0)$  is the short rate at time  $s$ ,  $h$  characterizes the cash flow at maturity, and  $\mathcal{F}_t$  is the information available at time  $t$ . This expectation is taken under the risk-neutral measure, reflecting the principle that the derivative price equals the present value of expected future payoffs when discounted at the risk-free rate.

Computing this expectation directly via Monte Carlo requires simulating the stochastic evolution of the entire forward curve from time  $t$  to time  $t + \tau$  along many paths, then averaging the discounted payoffs. The Feynman-Kac theorem transforms this stochastic problem into a deterministic PDE, avoiding the need for random simulation entirely.

**Theorem 3.5.1** (Multidimensional Discounted Feynman-Kac). *Consider a  $K$ -dimensional state vector  $X(t) \in \mathbb{R}^K$  evolving according to the stochastic differential equation:*

$$dX(t) = \mu(t, X(t))dt + \sigma(t, X(t))dW_t \quad (3.22)$$

where  $\mu(t, X(t)) \in \mathbb{R}^K$  is the drift vector,  $\sigma(t, X(t)) \in \mathbb{R}^{K \times N}$  is the diffusion matrix, and  $W_t \in \mathbb{R}^N$  is an  $N$ -dimensional Brownian motion. Let  $r(t, X(t))$  denote a discount rate that may depend on the state.

Define the pricing functional:

$$V(t, x) = \mathbb{E}^{t,x} \left[ e^{-\int_t^T r(s, X(s)) ds} h(T, X(T)) \right] \quad (3.23)$$

where  $\mathbb{E}^{t,x}[\cdot]$  denotes the conditional expectation given  $X(t) = x$ , and  $h(T, X(T))$  is the terminal payoff.

Then under general regularity conditions  $V$  satisfies the partial differential equation:

$$\frac{\partial V}{\partial t} + \mu' D_x V + \frac{1}{2} \sum_{n=1}^N \sigma'_n D_x^2 V \sigma_n - rV = 0 \quad (3.24)$$

subject to the terminal condition  $V(T, x) = h(T, x)$ , where  $D_x V$  is the gradient vector and  $D_x^2 V$  is the Hessian matrix of  $V$  with respect to  $x$ , and  $\sigma_n$  denotes the  $n$ -th column of the diffusion matrix.

The PDE (3.24) has an intuitive structure. The first term  $\partial V / \partial t$  captures the time evolution of the pricing functional. The second term  $\mu' D_x V$  represents the expected change in value due to the drift of the state variable. The third term  $\frac{1}{2} \sum_{n=1}^N \sigma'_n D_x^2 V \sigma_n$  captures the effect of volatility through second-order (convexity) terms. Finally, the term  $-rV$  discounts the value at the instantaneous rate.

To apply the Feynman-Kac theorem to the HJM framework, I must first discretize the forward curve over the tenor dimension. The forward curve  $f(t, \tau)$  is an infinite-dimensional object—a function mapping each tenor  $\tau \in [0, \bar{T}]$  to a forward rate. For computational purposes, I discretize this continuum by evaluating the forward curve at a finite grid of tenor points  $\{\tau_k\}_{k=1}^K$ . Define the discretized forward rates as:

$$f_k(t) = f(t, \tau_k) \quad \text{for } k = 1, \dots, K \quad (3.25)$$

and collect them into a state vector  $f(t) = (f_1(t), \dots, f_K(t))' \in \mathbb{R}^K$ . This vector represents the entire forward curve at time  $t$  through its values at the  $K$  discrete tenor points.

Under the Musiela parameterization presented earlier, the discretized forward curve evolves according to:

$$df(t) = \mu(t, f)dt + \sum_{n=1}^N \sigma_n(t, f)dW_n(t) \quad (3.26)$$

where the drift vector and volatility factors are defined as:

$$\mu(t, f) = (\mu(t, \tau_1, f), \dots, \mu(t, \tau_K, f))' \in \mathbb{R}^K \quad (3.27)$$

$$\sigma_n(t, f) = (\sigma_n(t, \tau_1, f), \dots, \sigma_n(t, \tau_K, f))' \in \mathbb{R}^K \quad \text{for } n = 1, \dots, N \quad (3.28)$$

and  $\{W_n(t)\}_{n=1}^N$  are the  $N$  independent Brownian motions driving term structure dynamics. This discretized system fits precisely into the framework of the multidimensional Feynman-Kac theorem with  $X(t) = f(t)$  and dimension  $K$ .

The no-arbitrage drift  $\mu(t, \tau_k, f)$  at each tenor point is computed from the Musiela condition as described in the previous section:

$$\mu(t, \tau_k, f) = \frac{\partial}{\partial \tau} f(t, \tau_k) + \sigma(t, \tau_k, f) \cdot \int_0^{\tau_k} \sigma(t, s, f) ds \quad (3.29)$$

where for the local volatility specification,  $\sigma(t, \tau_k, f) = \tilde{\sigma}(\tau_k) \cdot \min\{\sqrt{f_k(t)}, M\}$ . The discount rate in the Feynman-Kac theorem is the short rate  $r(t) = f(t, 0)$ .

Applying Theorem 1 directly yields the pricing PDE for interest rate derivatives in the HJM framework:

$$\frac{\partial V}{\partial t} + \mu(t, f)' D_f V + \frac{1}{2} \sum_{n=1}^N \sigma_n(t, f)' D_f^2 V \sigma_n(t, f) - r(t)V = 0 \quad (3.30)$$

subject to the terminal boundary condition:

$$V(t + \tau, f; \Xi) = h(f; \Xi) \quad (3.31)$$

where  $V(t, f; \Xi)$  is the time- $t$  value of the derivative when the forward curve is  $f$ , and the gradient vector

and Hessian matrix are defined as:

$$D_f V = \left( \frac{\partial V}{\partial f_1}, \dots, \frac{\partial V}{\partial f_K} \right)' \in \mathbb{R}^K \quad (3.32)$$

$$D_f^2 V = \left( \frac{\partial^2 V}{\partial f_k \partial f_\ell} \right)_{k,\ell=1}^K \in \mathbb{R}^{K \times K} \quad (3.33)$$

A crucial feature of this formulation is that the PDE (3.30) does *not* depend on the specific contract features  $\Xi$ . The contract-specific information enters only through the terminal boundary condition (3.31), which specifies the payoff function  $h(f; \Xi)$ . This separation has important practical implications: once I have trained a neural network to solve the PDE for one contract type, adapting to a different contract requires only changing the boundary condition in the loss function. The core PDE structure—the drift term, volatility terms, and discounting—remains unchanged across all interest rate derivatives priced under the same HJM model. This modularity makes the FINN approach highly flexible for pricing a wide range of instruments.

### 3.6 Application: Caplet Pricing

To demonstrate the Finance-Informed Neural Network approach, I apply it to pricing interest rate caplets. Caplets provide insurance against rising interest rates and are fundamental building blocks for more complex instruments such as interest rate caps (portfolios of caplets) and swaptions. The caplet pricing problem is particularly well-suited for demonstrating the FINN methodology because it is path-dependent in multiple ways: the payoff depends on the future LIBOR rate, which itself depends on the future forward curve, and the entire cash flow must be discounted along the stochastic path of short rates. This path dependence makes Monte Carlo simulation computationally expensive, yet the payoff structure is straightforward enough to allow clear interpretation of results.

A caplet is a call option on a future LIBOR rate. To define the contract precisely, I first introduce the LIBOR rate and then specify the caplet payoff.

The LIBOR (London Interbank Offered Rate) is an annualized interest rate for borrowing between time  $\tau_1$  and time  $\tau_2$ .<sup>3</sup> At time  $t$ , the forward LIBOR rate starting at  $\tau_1$  and ending at  $\tau_2$  is defined implicitly through the relationship between bond prices:

$$L(t, \tau_1, \tau_2) = \frac{1}{\delta} \left( \frac{P(t, \tau_1)}{P(t, \tau_2)} - 1 \right) \quad (3.34)$$

where  $\delta = \tau_2 - \tau_1$  is the accrual period (typically 3 months or 6 months for LIBOR), and  $P(t, \tau)$  is the time- $t$  price of a zero-coupon bond maturing at time  $t + \tau$  with face value normalized to one. Under the HJM framework, bond prices are determined by integrating the forward curve:

$$P(t, \tau) = \exp \left( - \int_0^\tau f(t, s) ds \right) \quad (3.35)$$

Thus the LIBOR rate depends on the entire forward curve from 0 to  $\tau_2$  through the bond price ratio,

---

<sup>3</sup>While LIBOR has been phased out in recent years in favor of risk-free rates such as SOFR, the caplet pricing problem remains relevant for understanding interest rate derivatives more generally.

since both  $P(t, \tau_1)$  and  $P(t, \tau_2)$  require integrating forward rates from zero to their respective maturities.

A caplet is a call option on the LIBOR rate  $L(t, \tau_1, \tau_2)$  with strike price  $L_E$ . The holder of the caplet receives a payoff at time  $t + \tau_2$  (the end of the accrual period) equal to:

$$\text{Payoff at time } t + \tau_2 = \delta \cdot \max\{L(t + \tau_1, \tau_1, \tau_2) - L_E, 0\} \quad (3.36)$$

where the LIBOR rate is fixed at the start of the accrual period ( $t + \tau_1$ ) and the payoff is received at the end ( $t + \tau_2$ ). The factor  $\delta$  converts the annualized rate difference into a dollar amount over the accrual period.

For computational convenience, I adopt the market convention of valuing the caplet at the settlement date  $t + \tau_1$  rather than the payment date  $t + \tau_2$ . Discounting the payoff back one period yields the time- $(t + \tau_1)$  value:

$$V(t + \tau_1, 0, f) = \delta P(t + \tau_1, \delta) \max\{L(t + \tau_1, 0, \delta) - L_E, 0\} \quad (3.37)$$

where I have simplified notation by setting  $\tau_1 = 0$  relative to the valuation date and  $\tau_2 = \delta$ . The contract features are summarized as  $\Xi = (\tau, \delta, L_E)$  where  $\tau$  is the time from today until the settlement date  $t + \tau_1$ ,  $\delta$  is the accrual period, and  $L_E$  is the strike price.

### 3.6.1 Neural Network Strategy

Having derived the pricing PDE (3.30), the next challenge is to solve it numerically. In principle, the PDE requires computing the pricing functional  $V(t, f; \Xi)$  and its derivatives— $\partial V / \partial t$ , the gradient  $D_f V$ , and the Hessian  $D_f^2 V$ —at many points in the high-dimensional state space  $(t, f) \in \mathbb{R} \times \mathbb{R}^K$ . However, a key simplification arises from the Musiela parameterization: since the contract features  $\Xi$  already include the time-to-maturity  $\tau$ , and the pricing problem can be solved backward from the settlement date, I can eliminate explicit dependence on calendar time  $t$ . The pricing functional becomes  $V(\tau, f; \Xi)$  where

$\tau$  represents the remaining time until settlement. Under the Musiela parameterization, since  $\tau = T - t$  where  $T$  is the fixed settlement date, the chain rule yields  $\frac{\partial V}{\partial t} = -\frac{\partial V}{\partial \tau}$ . This transforms the time derivative into a derivative with respect to time-to-maturity, allowing the neural network to be parameterized directly as a function of  $(\tau, f, \Xi)$  rather than requiring separate tracking of calendar time.

Traditional finite difference methods for solving the resulting PDE become prohibitively expensive in high dimensions due to the curse of dimensionality: discretizing each dimension of the forward curve on a grid leads to exponential growth in the number of grid points. Neural networks offer a solution. A neural network is a smooth, differentiable function that can approximate complex nonlinear mappings. I parameterize the pricing function as a neural network  $V_{\Theta}(f; \Xi)$  where  $\Theta$  represents the collection of all weights and biases in the network. The network takes as inputs the discretized forward curve  $f = (f_1, \dots, f_K)$  and contract features  $\Xi$  (which include time-to-maturity  $\tau$ ), and outputs a scalar price estimate.

A key computational advantage lies in *automatic differentiation*. Modern deep learning frameworks such as JAX, PyTorch, and TensorFlow implement automatic differentiation systems that can compute derivatives of any function constructed through their operations—including arbitrarily complex neural networks. Critically, these derivatives are computed exactly (up to floating-point precision), not through finite difference approximations. When I evaluate the neural network  $V_{\Theta}(f; \Xi)$ , the automatic differentiation system tracks all operations and can immediately compute the first- and second-order derivatives needed for evaluating the PDE governing the pricing functional. These derivative computations are efficient and scale easily with dimension  $K$ , unlike grid-based finite difference methods. An additional practical benefit is that the Greeks—sensitivities of the option price to market parameters—come essentially for free. Since the pricing PDE itself contains  $\partial V / \partial \tau$  (theta, time decay) and  $D_f V$  (the gradient yielding curve deltas, sensitivities to each forward rate  $f_k$ ), these quantities are computed automatically when evaluating the PDE residual during training. Thus, evaluating pricing errors simultaneously delivers the Greeks critical for risk management and hedging strategies at zero marginal cost.

With automatic differentiation providing all necessary derivatives, I can directly evaluate how well the

neural network satisfies the PDE at any point  $(\tau, f)$ . The PDE residual at a point is:

$$-\frac{\partial V_{\Theta}}{\partial \tau} + \mu(t, f)' D_f V_{\Theta} + \frac{1}{2} \sum_{n=1}^N \sigma_n(t, f)' D_f^2 V_{\Theta} \sigma_n(t, f) - r(t) V_{\Theta} \quad (3.38)$$

### 3.6.2 FINN Architecture and Training

For the remainder of this section, I adopt the notation  $\mathcal{X}$  to denote the inputs to the neural network and  $\mathcal{Y}$  to denote the outputs. The inputs are  $\mathcal{X} = (f_1, \dots, f_K, S, \Xi) \in \mathbb{R}^{K+9}$  comprising the discretized forward curve, Svensson parameters  $S = (\beta_0, \beta_1, \beta_2, \beta_3, \tau_1, \tau_2) \in \mathbb{R}^6$ , and contract features  $\Xi = (\tau, \delta, L_E) \in \mathbb{R}^3$  where  $\tau$  is the time until the start of the accrual period,  $\delta$  is the accrual length, and  $L_E$  is the strike price. Note that  $\tau_2 = \tau + \delta$  is redundant and not included as a separate input. The output is the scalar caplet price  $\mathcal{Y} = V(\tau, f; \Xi)$ .

To improve training stability and convergence, I normalize inputs within the network architecture itself—crucially, as the first layer before any nonlinear transformations. This normalization is implemented as a differentiable operation within the computational graph, meaning automatic differentiation seamlessly handles all derivative computations through the normalization without requiring manual adjustment. The time-related contract features  $\tau_1$  and  $\delta$  are normalized by dividing by the maximum maturity  $\tau_{\max}$  in the data. The Svensson parameters are transformed to z-scores:

$$\tilde{\beta}_i = \frac{\beta_i - \mu_{\beta_i}}{\sigma_{\beta_i}} \quad (3.39)$$

where  $\mu_{\beta_i}$  and  $\sigma_{\beta_i}$  are the empirical mean and standard deviation of parameter  $\beta_i$  computed across the historical dataset. The forward rates themselves are left unnormalized, as their scale is economically meaningful and directly enters the pricing formula.

The FINN architecture consists of three hidden layers, each with 500 neurons. I apply the sigmoid linear unit (SiLU) activation function  $\text{silu}(x) = x \times \sigma(x)$  where  $\sigma(x) = \frac{1}{1+\exp(-x)}$  is the sigmoid function

to the hidden layers, which provides smooth, differentiable nonlinearity necessary for backpropagation through the financial equilibrium conditions. Since prices must be nonnegative, I apply the softplus activation function  $\text{softplus}(x) = \log(1 + e^x)$  to the output layer, which smoothly enforces the non-negativity constraint while remaining differentiable everywhere.

To accelerate training, I precompute several quantities that remain constant across training iterations. First, I enumerate all admissible tenor-accrual pairs  $(\tau, \delta)$ , storing these as indices for rapid sampling during batch generation. Second, since all forward rates and volatility evaluations occur on the fixed tenor grid  $\{\tau_k\}_{k=1}^K$ , I precompute the trapezoidal integration matrix  $C \in \mathbb{R}^{K \times K}$  where element  $C_{ij}$  gives the weight for integrating from grid point  $j$  to grid point  $i$ . This vectorizes all integral computations—including bond price calculations  $P(t, \tau) = \exp(-\int_0^\tau f(t, s)ds)$  and drift integrals—into matrix-vector products. Third, I precompute and store the Chebyshev polynomial coefficients for the volatility functions  $\sigma_n(\tau)$  from the PCA decomposition. For the constant volatility specification, I additionally precompute the volatility integral  $\int_0^\tau \sigma(t, s)ds$  needed in the drift term, storing these as vectors that can be directly used in the loss function evaluation. These precomputations transform computationally expensive operations into simple lookups and matrix multiplications, substantially reducing per-iteration training time.

Unlike traditional supervised learning where training data comes from a pre-existing dataset, the FINN approach generates its own training data by sampling from the historical forward curve distribution. Each training batch is constructed by randomly sampling:

- A tenor-accrual pair with  $\tau \in [0, 5]$  and  $\delta \in [1/3, 3/4]$  (covering typical caplet accrual periods and maturities)
- A forward curve  $f$  and its corresponding Svensson parameters from the filtered historical dataset
- A strike price  $L_E$  from a Chebyshev grid over  $[0, 0.07]$  (covering the range of observed forward rates)

This sampling strategy ensures the network trains on a diverse set of market conditions, contract specifications, and moneyness levels. The training data is regenerated every epoch, preventing overfitting to any particular set of forward curves and encouraging the network to learn the underlying PDE structure rather than memorizing specific curve realizations.

To evaluate the boundary condition, for any historically observed forward curve  $f$  sampled from the data, I simply set the time-to-maturity  $\tau = 0$  in the contract features and evaluate what the payoff would be if that curve represented the state at expiry. That is, I treat each sampled historical curve as if it were the forward curve at the settlement date and compute the corresponding caplet payoff. The loss function then penalizes deviations between the network's prediction  $V_{\Theta}(\tau = 0, f; \Xi)$  and this analytically computed payoff. This approach leverages the rich historical variation in forward curve shapes to train the boundary condition, while the PDE loss teaches the network how values at earlier times  $\tau > 0$  relate to these terminal payoffs through the no-arbitrage dynamics. Importantly, this strategy completely avoids the computational expense of path simulation—all training occurs on the cross-section of historical curves, not their time-series evolution.

To further enrich the disciplining of the pricing functional, I exploit the special case of zero-strike caplets, which admit a closed-form analytical expression that holds at all times, not just at maturity. For a caplet with strike  $L_E = 0$ , there is no optionality—the payoff is deterministically equal to the LIBOR rate. Standard no-arbitrage arguments yield the closed-form expression:

$$V(\tau_1, f; L_E = 0) = P(\tau_1) - P(\tau_1 + \delta) = \exp\left(-\int_0^{\tau_1} f(s)ds\right) - \exp\left(-\int_0^{\tau_1 + \delta} f(s)ds\right) \quad (3.40)$$

Crucially, this closed-form expression holds for any  $\tau_1 \geq 0$ , not just at the boundary. The absence of optionality means the pricing PDE reduces to a deterministic bond pricing formula, which can be evaluated exactly from the current forward curve. During training, I augment the loss function by relabeling a fraction of the sampled data points to have strike  $L_E = 0$  at various times-to-maturity  $\tau_1$  and penalizing deviations between the network's zero-strike prediction and this analytical formula. This pro-

vides the network with exact supervisory signals across the full domain—both in the time dimension and across the space of forward curve shapes—serving as a powerful regularization that anchors the pricing functional to known analytical values. The zero-strike case effectively provides thousands of additional training points where the true solution is known exactly, complementing the PDE residual minimization and boundary condition training for positive strikes, thereby improving overall accuracy and accelerating convergence.

The network is trained using the Adam optimizer (Kingma and Ba 2014) with weight decay regularization ( $\lambda = 10^{-5}$ ) in a three-regime curriculum learning schedule. Each regime uses progressively smaller learning rates to refine the solution:

- Regime 1: 15,000 epochs with learning rate  $10^{-4}$ , batch size 100, 10 batches per epoch
- Regime 2: 5,000 epochs with learning rate  $10^{-5}$ , batch size 100, 10 batches per epoch
- Regime 3: 2,500 epochs with learning rate  $10^{-6}$ , batch size 500, 2 batches per epoch

This curriculum learning approach starts with larger learning rates to explore the solution space broadly, then progressively refines with smaller learning rates and larger batch sizes. Crucially, the training data is redrawn after each epoch, ensuring that the network learns from fresh forward curve realizations rather than overfitting to a single sample.

A key computational challenge in evaluating the PDE residual is computing the second-order term, which involves the Hessian matrix  $D_f^2 V \in \mathbb{R}^{K \times K}$ . I employ a directional derivative trick to avoid materializing the full Hessian. Using the chain rule identity, the quadratic form can be rewritten as:

$$\sigma'_n D_f^2 V \sigma_n = D_f(D_f V \cdot \sigma_n) \cdot \sigma_n \quad (3.41)$$

The algorithm proceeds as follows: first compute the gradient  $D_f V$  once. Then for each factor  $n$ , form the scalar directional derivative  $s = D_f V \cdot \sigma_n$ , compute its gradient  $D_f s$  (which gives the second deriva-

tive in the  $\sigma_n$  direction), and dot with  $\sigma_n$  again to obtain  $\sigma'_n D_f^2 V \sigma_n$ . This approach reduces memory and computational requirements, making it feasible to evaluate the PDE residual efficiently even for large discretizations  $K$ .

Combining these elements, the complete training objective becomes clear. If  $V_\Theta$  were the true solution, this residual would equal zero everywhere. I train the neural network to minimize the squared PDE residual across a collection of randomly sampled points, combined with penalties for violating both the terminal boundary condition and the zero-strike analytical formula. The Finance-Informed loss function comprises three components, averaged over number of contracts  $N_{\text{batch}}$ :

$$\begin{aligned} \mathcal{L}_\Theta = \frac{1}{N_{\text{batch}}} \sum_{n=1}^{N_{\text{batch}}} \left( \left( -\frac{\partial V_\Theta}{\partial \tau_1} + \mu' D_f V_\Theta + \frac{1}{2} \sum_{n=1}^N \sigma'_n D_f^2 V_\Theta \sigma_n - r V_\Theta \right)^2 \right. \\ \left. + (V_\Theta(\tau_1 = 0, f; \Xi) - \delta P(\delta) \max\{L(0, \delta) - L_E, 0\})^2 \right. \\ \left. + (V_\Theta(\tau_1, f; L_E = 0) - (P(\tau_1) - P(\tau_1 + \delta)))^2 \right) \quad (3.42) \end{aligned}$$

where the first term penalizes violations of the pricing PDE across sampled interior points  $(\tau_1, f)$  with  $\tau_1 > 0$ , the second term enforces the terminal boundary condition at settlement  $(\tau_1 = 0)$  for positive-strike caplets, and the third term anchors the network to the analytical zero-strike formula across all times-to-maturity. The neural network parameters are updated through standard gradient descent (specifically, the Adam optimizer) by solving:

$$\Theta^* = \arg \min_{\Theta} \mathcal{L}_\Theta \quad (3.43)$$

This approach—using automatic differentiation to enforce PDE structure through a loss function—allows the neural network to learn the pricing function without ever simulating Monte Carlo paths. The method scales efficiently to high-dimensional state spaces and handles path-dependent payoffs through the PDE framework rather than through stochastic simulation.

Table 3.1 reports the hardware used in training the model. Since the FINN generates its own training

data rather than requiring large pre-existing datasets, the memory bottleneck is the network size rather than dataset size. This makes consumer-grade hardware with modest GPU memory (8GB VRAM) entirely sufficient, eliminating the need for specialized high-performance computing infrastructure typical of large-scale machine learning applications.

Table 3.1: Computational Environment

| Hardware  |                                           |
|-----------|-------------------------------------------|
| Processor | 12th Gen Intel i9-12900KF (24) @ 5.100GHz |
| GPU       | NVIDIA GeForce RTX 3080 with 8GB of VRAM  |
| RAM       | 32 GB                                     |

To assess how solution accuracy and computational cost scale with the dimension of the discretized forward curve, I train separate models for eight different discretizations, varying the parameter  $K \in \{10, 25, 35, 50, 75, 100, 125, 150\}$ . Remarkably, training time remains approximately one hour across all discretization levels, demonstrating that the FINN approach scales gracefully with dimension—a stark contrast to traditional finite difference or Monte Carlo methods where computational cost grows exponentially or linearly (respectively) with  $K$ . This dimensional robustness arises from the efficiency of automatic differentiation and the fact that neural network evaluation cost grows only modestly with input dimension.

### 3.6.3 Results

I validate the FINN approach by comparing caplet prices to Monte Carlo benchmarks across all eight discretization levels  $K \in \{10, 25, 35, 50, 75, 100, 125, 150\}$ . All results employ the local volatility specification described above, where volatility scales with the square root of the forward rate level. For each model, I generate a test set of 1,000 randomly sampled forward curves with varying caplet contract specifications (strikes, tenors, and accrual periods) and compute both FINN prices and Monte Carlo prices using 10,000 simulated paths per contract. The Monte Carlo benchmark likewise uses the same local volatility structure, ensuring a consistent comparison. The results demonstrate that FINNs achieve comparable pricing accuracy to Monte Carlo simulation while delivering transformative computational

speedups.

Figure 3.1 quantifies pricing accuracy by plotting the mean absolute error between FINN and Monte Carlo prices as a function of  $K$ . The error profile does not decrease monotonically with  $K$ —instead, it exhibits a non-monotonic pattern with local minima around  $K = 50$  and  $K = 100$ , and local peaks around  $K = 10$  and  $K = 75$ . All errors remain below 0.001 (0.1 cents on a dollar-denominated contract), with most discretizations achieving errors between 0.0004 and 0.0007. The finest discretizations ( $K = 100$  and  $K = 150$ ) both achieve errors around 0.0004, representing approximately 0.04 cents per dollar of contract value. This non-monotonicity suggests that discretization level interacts with the training dynamics and neural network capacity in complex ways, rather than finer grids uniformly improving accuracy. For practical purposes, all discretization levels achieve sufficient accuracy for trading applications, with  $K = 100$  and  $K = 150$  providing the tightest error bounds.




Figure 3.1: Mean Absolute Pricing Error vs. Discretization Level (1000-contract test set). All discretizations maintain error below 0.001, well within acceptable bounds for trading applications. The finest discretizations ( $K = 100$  and  $K = 150$ ) achieve the tightest error bounds around 0.0004. The non-monotonicity suggests complex interactions between discretization level, training dynamics, and network capacity.

An important caveat: Monte Carlo prices themselves are not exact benchmarks but rather approximations subject to their own sources of error. The Monte Carlo implementation discretizes the forward curve into  $K$  tenor points (the same discretization used by the FINN), discretizes time into finite steps (Euler-Maruyama Markov chain), and performs repeated numerical integrations to produce drift terms.

Crucially, when computing the drift term, Monte Carlo must approximate the tenor derivative  $\partial f / \partial \tau$  using finite differences on the discretized grid, introducing additional discretization error. By contrast, the FINN leverages the analytical Svensson formula for  $\partial f / \partial \tau$ , avoiding this source of approximation entirely. Additionally, Monte Carlo prices are subject to sampling error despite using 10,000 paths per contract. The reported errors therefore reflect the combined approximation errors from both methods rather than pure FINN error relative to a known analytical solution. In principle, the FINN could be closer to the true price than the Monte Carlo benchmark, particularly if the neural network captures the smooth PDE solution more accurately than discrete-time simulation with finite paths. The sub-0.001 agreement between methods suggests both approaches achieve high accuracy, but neither should be viewed as providing ground truth.




Figure 3.2: Final Training Loss Across Discretizations. Both PDE residual (solid black) and boundary condition violation (dashed blue) remain stable across all values of  $K$ , with losses on the order of  $10^{-7}$  to  $10^{-8}$ . The consistency of loss magnitudes demonstrates that the FINN training procedure scales robustly with dimension.

Figure 3.2 displays the final PDE loss (solid black line) and boundary condition (BC) loss (dashed blue line) achieved after training for each discretization level. Both loss components remain remarkably stable across  $K$ , with PDE loss fluctuating around  $5 \times 10^{-8}$  and BC loss varying between approximately  $5 \times 10^{-8}$  and  $4 \times 10^{-7}$ . The y-axis uses a logarithmic scale, highlighting that final losses are uniformly small (on the order of  $10^{-7}$  to  $10^{-8}$ ), indicating high-precision satisfaction of both the pricing PDE and terminal boundary conditions across all discretization levels. The consistency of these loss values con-

firmly that the FINN training procedure scales robustly with dimension—there is no evidence of training degradation at higher  $K$ .

The most compelling advantage of the FINN approach lies in computational speed. Figure 3.3 presents comprehensive comparisons of FINN and Monte Carlo evaluation times. The top panel of Figure 3.3 shows pricing time in seconds per contract on a linear scale: Monte Carlo time (solid black line) grows from approximately 1 second at  $K = 10$  to nearly 10 seconds at  $K = 150$ , while FINN time (dashed blue line) remains essentially flat near zero across all discretizations. The middle panel displays the same data on a log scale, revealing that FINN evaluation takes approximately  $10^{-6}$  to  $10^{-5}$  seconds per contract (a few microseconds), while Monte Carlo ranges from 1 to 10 seconds. The bottom panel shows the speedup ratio (MC time / FINN time), which grows from approximately 300,000 at  $K = 10$  to over 4.5 million at  $K = 150$ —representing a *hundred-thousand to multi-million-fold speedup*. Table 3.2 provides the precise timing measurements and mean absolute errors for each discretization level.

This is the headline result: FINNs price caplets roughly *300,000 to 4,500,000 times faster* than Monte Carlo simulation, with the advantage increasing as dimension grows. Once trained, the neural network evaluates in a few microseconds regardless of the forward curve dimensionality, while Monte Carlo evaluation time grows linearly with  $K$ . This dramatic speed advantage makes FINNs transformative for applications requiring rapid repricing of large portfolios, such as real-time risk management, high-frequency trading, or iterative calibration procedures.




Figure 3.3: Pricing Speed: FINNs vs. Monte Carlo (1000-contract test set). Top panel: Monte Carlo time (solid) grows linearly from 1 to 10 seconds per contract as  $K$  increases, while FINN time (dashed) remains near zero. Middle panel (log scale): FINN evaluation takes  $\sim 10^{-6}$  to  $10^{-5}$  seconds (a few microseconds), while MC takes 1 to 10 seconds. Bottom panel: Speedup ratio grows from approximately 300,000 $\times$  at  $K = 10$  to over 4.5 million $\times$  at  $K = 150$ , representing a multi-million-fold computational advantage that increases with dimension.

| $K$ | MC Time | FINN Time             | Speedup Multiple | MAE                   |
|-----|---------|-----------------------|------------------|-----------------------|
| 10  | 0.96    | $3.10 \times 10^{-6}$ | 311,366          | $8.73 \times 10^{-4}$ |
| 25  | 1.18    | $4.25 \times 10^{-6}$ | 276,943          | $6.25 \times 10^{-4}$ |
| 35  | 1.40    | $5.45 \times 10^{-6}$ | 257,440          | $5.69 \times 10^{-4}$ |
| 50  | 1.92    | $5.91 \times 10^{-6}$ | 324,690          | $5.17 \times 10^{-4}$ |
| 75  | 3.49    | $2.45 \times 10^{-6}$ | 1,421,537        | $9.42 \times 10^{-4}$ |
| 100 | 5.44    | $2.16 \times 10^{-6}$ | 2,522,035        | $3.75 \times 10^{-4}$ |
| 125 | 8.22    | $2.53 \times 10^{-6}$ | 3,245,742        | $7.98 \times 10^{-4}$ |
| 150 | 10.08   | $2.23 \times 10^{-6}$ | 4,522,638        | $4.03 \times 10^{-4}$ |

Table 3.2: Computational performance and accuracy across discretization phases. MC Time and FINN Time are in seconds per contract. Speedup Multiple is the ratio of MC Time to FINN Time. MAE (Mean Absolute Error) measures the difference between FINN and Monte Carlo prices.

Figure 3.4 plots FINN prices against Monte Carlo prices for all eight discretization levels. Each panel corresponds to a different value of  $K$ , with points colored by strike price—purple indicates near-zero strikes (high caplet prices near the top-right of each panel), while yellow indicates high strikes near 7% (lower prices, deep out-of-the-money contracts near the origin). The scatter plots reveal tight clustering along the 45-degree line across all discretizations, demonstrating strong agreement between FINN and Monte Carlo prices. Notably, the near-zero strike contracts (purple points with highest prices) exhibit particularly tight agreement, reflecting the zero-strike analytical anchoring used during training. However, higher-strike contracts (yellow/green points at lower prices, farther from the zero-strike anchor) show slightly more scatter, particularly visible in panels with coarser discretizations. This suggests that the FINN struggles more with contracts farther from the analytical anchor point, pointing to potential future work: anchoring from both the zero-strike and high-strike extremes could improve accuracy across the full moneyness spectrum.




Figure 3.4: FINN vs. Monte Carlo Prices Across Discretizations (1000-contract test set). Each panel shows test contracts for a different discretization level  $K$ , with points colored by strike price  $L_E$ . Points cluster tightly along the 45-degree line, with near-zero strike contracts (purple) showing the tightest agreement due to the analytical zero-strike anchoring. Higher-strike contracts (yellow/green) exhibit slightly more scatter, suggesting the FINN performs best near the analytical anchor and pointing to future work on dual anchoring from both extremes.

## 3.7 Conclusion

This paper demonstrates a fundamental computational breakthrough for pricing path-dependent interest rate derivatives under the Heath-Jarrow-Morton framework. The HJM model offers unparalleled generality: it specifies arbitrage-free dynamics for the entire forward curve simultaneously, nesting earlier short-rate models as special cases and providing a theoretically rigorous foundation for derivatives pricing. However, this generality has historically come at a severe computational cost. The forward curve is infinite-dimensional, and pricing path-dependent contracts traditionally requires Monte Carlo simulation of thousands or millions of stochastic paths—a computational burden that grows linearly with the dimensionality of the discretized state space.

I circumvent Monte Carlo simulation entirely by invoking the Feynman-Kac theorem, which establishes that stochastic expectations can be characterized as solutions to deterministic partial differential equations. Rather than simulating random forward curve paths to estimate expected payoffs, I solve the PDE governing the pricing functional directly. This transformation eliminates the need for path generation, replacing stochastic simulation with a deterministic PDE plus boundary value problem. However, this substitution merely shifts the computational challenge: traditional finite difference and finite element methods for solving high-dimensional PDEs suffer from the curse of dimensionality, with memory and computational requirements growing exponentially in the number of state variables.

The solution lies in Finance-Informed Neural Networks (FINNs)—deep learning models trained to satisfy the pricing PDE by embedding the differential equation directly into the loss function. The neural network parameterizes the pricing functional, and automatic differentiation computes the exact derivatives needed to evaluate PDE residuals at any point in the state space. Training minimizes violations of the governing PDE across sampled forward curve realizations, combined with penalties for violating terminal boundary conditions and deviations from analytical zero-strike solutions. This approach leverages three key advantages of neural networks: universal approximation capability for complex nonlinear functions, smooth differentiability enabling automatic computation of all required derivatives, and graceful scaling

with input dimension.

Crucially, FINN evaluation cost does not grow with the size of the state space. Once trained, the neural network prices derivatives nearly instantaneously regardless of whether the discretized forward curve contains ten tenor points or one hundred and fifty. Monte Carlo methods, by contrast, exhibit computational cost that grows linearly with discretization level—each additional state variable requires simulating more data and more computational expense in the form of integrals. This dimensional robustness makes FINNs transformative for pricing derivatives in high-dimensional settings where traditional methods become prohibitively expensive.

An additional practical advantage emerges from the structure of the pricing PDE itself: the major Greeks—theta (time decay) and curve deltas (sensitivities to each forward rate)—appear directly as terms in the differential equation. Since automatic differentiation computes these quantities when evaluating the PDE residual during training, they are obtained at zero marginal cost once the network is trained. Other Greeks, such as gamma (convexity) and vega (volatility sensitivity), require computing additional derivatives but remain computationally cheap via the same automatic differentiation framework. This stands in stark contrast to Monte Carlo methods, which cannot provide Greeks without complete re-simulation. To compute a single delta via Monte Carlo requires perturbing the contract parameter and running thousands of paths anew—and because even small parameter perturbations can significantly alter simulated prices through the stochastic dynamics, each Greek calculation demands a fresh Monte Carlo simulation. For a portfolio requiring hundreds of prices and thousands of sensitivities, this quickly becomes computationally prohibitive. Once the FINN is trained, however, all prices and all Greeks are available functionally and instantaneously through simple forward and backward passes of the neural network, with automatic differentiation delivering exact derivatives at negligible cost. For practitioners managing large derivatives portfolios, this represents a fundamental improvement over Monte Carlo: the ability to compute thousands of prices and their complete associated risk profiles nearly instantaneously, without any re-simulation.

The empirical results validate this approach decisively. Tested on 1,000 randomly sampled caplet con-

tracts across eight discretization levels ( $K \in \{10, 25, 35, 50, 75, 100, 125, 150\}$ ), FINNs achieve pricing accuracy within 0.04 to 0.07 cents per dollar of contract value compared to Monte Carlo benchmarks. These errors are well within acceptable bounds for trading applications, particularly considering that the Monte Carlo benchmark itself is subject to multiple sources of approximation error: discretization of the forward curve into  $K$  nodes, discrete time-stepping via Euler-Maruyama, sampling error despite using 10,000 paths, and crucially, approximation of the tenor derivative  $\partial f / \partial \tau$  via finite differences on the discretized grid. The FINN, by contrast, uses the analytical Svensson formula for this derivative, avoiding this source of discretization error. Given that both methods involve approximations, the sub-0.1 cent agreement between FINN and Monte Carlo prices demonstrates that FINNs achieve pricing accuracy comparable to—and possibly exceeding—traditional simulation methods. The computational advantage is dramatic: FINNs price caplets 300,000 to 4,500,000 times faster than Monte Carlo simulation, with speedups increasing as dimension grows. Once trained on consumer-grade hardware (8GB GPU), evaluation takes only a few microseconds per contract regardless of discretization level. Monte Carlo evaluation time, by contrast, grows linearly with  $K$ , reaching nearly 10 seconds per contract at  $K = 150$ . This scaling behavior makes FINNs particularly attractive for real-time risk management, high-frequency trading, iterative model calibration, and any application requiring rapid repricing of large derivatives portfolios under varying market conditions.

Beyond the immediate application to caplet pricing, the FINN methodology generalizes naturally to other path-dependent interest rate derivatives—caps, floors, swaptions, callable bonds—all of which can be priced by modifying only the boundary condition in the loss function while retaining the same core PDE structure. The framework already accommodates local volatility (all results in this paper use volatility that scales with the square root of forward rate levels), and extends equally to more sophisticated stochastic volatility models and jump-diffusion processes, with state-dependent coefficients entering seamlessly through automatic differentiation. This flexibility, combined with the dramatic computational speedups and essentially free Greeks, positions Finance-Informed Neural Networks as a powerful new tool for derivatives pricing in high-dimensional continuous-time models.

## Bibliography

- Abbott, Brant et al. (2019). “Education policy and intergenerational transfers in equilibrium”. In: *Journal of Political Economy* 127.6, pp. 2569–2624.
- Arnoud, Antoine et al. (2021). “The evolution of us firms’ retirement plan offerings: Evidence from a new panel data set”. In: *NBER Working Paper*.
- Auerbach, Alan J and Laurence J Kotlikoff (1987). “Evaluating fiscal policy with a dynamic simulation model”. In: *The American Economic Review* 77.2, pp. 49–55.
- Azinovic, Marlon, Luca Gaegauf, and Simon Scheidegger (2019). “Deep equilibrium nets”. In: *Available at SSRN 3393482*.
- Azinovic, Marlon and Jan Zemlicka (2024). *Intergenerational consequences of rare disasters*.
- Beck, Christian et al. (2021). “Solving the Kolmogorov PDE by means of deep learning”. In: *Journal of Scientific Computing* 88.3, p. 73.
- Bhargava, Saurabh and Lynn Conell-Price (2022). “Serenity Now, Save Later? Evidence on Retirement Savings Puzzles from a 401 (k) Field Experiment”. In.
- Black, Fischer (1976). “The pricing of commodity contracts”. In: *Journal of financial economics* 3.1-2, pp. 167–179.
- Black, Fischer and Myron Scholes (1973). “The pricing of options and corporate liabilities”. In: *Journal of political economy* 81.3, pp. 637–654.
- Bleemer, Zachary and Aashish Mehta (Apr. 2022). “Will Studying Economics Make You Rich? A Regression Discontinuity Analysis of the Returns to College Major”. In: *American Economic Journal*:

- Applied Economics* 14.2, pp. 1–22. DOI: [10.1257/app.20200447](https://doi.org/10.1257/app.20200447). URL: <https://www.aeaweb.org/articles?id=10.1257/app.20200447>.
- Bloomfield, Adam et al. (July 2025). *How Do Tax Incentives Influence Employer Decisions to Offer Retirement Benefits?* Working Paper 34043. National Bureau of Economic Research. DOI: [10.3386/w34043](https://doi.org/10.3386/w34043). URL: <http://www.nber.org/papers/w34043>.
- Board of Governors of the Federal Reserve System (2019). *Survey of Consumer Finances, 2019*. Tech. rep. Washington, DC: Board of Governors of the Federal Reserve System. URL: <https://www.federalreserve.gov/econres/scfindex.htm>.
- Board of Governors of the Federal Reserve System [US] (Aug. 2022). “Student Loans Owned and Securitized (SLOAS)”. In: *FRED, Federal Reserve Bank of St. Louis*. URL: <https://fred.stlouisfed.org/graph/?g=R5m4>.
- Boutros, Michael, Nuno Clara, and Francisco Gomes (2024). “Borrow now, pay even later: A quantitative analysis of student debt payment plans”. In: *Journal of Financial Economics* 159, p. 103898.
- Brace, Alan, Dariusz Gałtarek, and Marek Musiela (1997). “The market model of interest rate dynamics”. In: *Mathematical finance* 7.2, pp. 127–155.
- Campbell, John Y (2016). “Restoring rational choice: The challenge of consumer financial regulation”. In: *American Economic Review* 106.5, pp. 1–30.
- Cao, Jay, Jacky Chen, Soroush Farghadani, et al. (2023). “Gamma and vega hedging using deep distributional reinforcement learning”. In: *Frontiers in Artificial Intelligence* 6, p. 1129370.
- Cao, Jay, Jacky Chen, and John Hull (2020). “A neural network approach to understanding implied volatility movements”. In: *Quantitative Finance* 20.9, pp. 1405–1413.
- Cao, Jay, Jacky Chen, John Hull, and Zissis Poulos (2021). “Deep learning for exotic option valuation”. In: *arXiv preprint arXiv:2103.12551*.
- Cass, David and Karl Shell (1983). “Do sunspots matter?”. In: *Journal of political economy* 91.2, pp. 193–227.
- Catherine, Sylvain and Constantine Yannelis (2023). “The distributional effects of student loan forgiveness”. In: *Journal of Financial Economics* 147.2, pp. 297–316.

- Chakrabarti, Rajashri et al. (2020). “Tuition, debt, and human capital”. In: *FRB of New York Staff Report* 912.
- Chen, Jacky et al. (2023). *A Variational Autoencoder Approach to Conditional Generation of Possible Future Volatility Surfaces*.
- Cheyette, Oren (2001). “Markov representation of the Heath-Jarrow-Morton model”. In: *Available at SSRN* 6073.
- Citanna, Alessandro and Paolo Siconolfi (2007). “Short-memory equilibrium in stochastic overlapping generations economies”. In: *Journal of Economic Theory* 134.1, pp. 448–469.
- (2010). “Recursive Equilibrium in Stochastic Overlapping Generations Economies”. In: *Econometrica* 78.1, pp. 309–347.
- Coimbra, Nuno et al. (2023). “Asset Pricing and Risk-Sharing Implications of Alternative Pension Plan Systems”. In: *The Journal of Finance*.
- Constantinides, George M, John B Donaldson, and Rajnish Mehra (2002). “Junior can’t borrow: A new perspective on the equity premium puzzle”. In: *The Quarterly Journal of Economics* 117.1, pp. 269–296.
- Cox, John C, Jonathan E Ingersoll, Stephen A Ross, et al. (1985). “A theory of the term structure of interest rates”. In: *Econometrica* 53.2, pp. 385–407.
- Dammon, Robert M, Chester S Spatt, and Harold H Zhang (2004). “Optimal asset location and allocation with taxable and tax-deferred investing”. In: *The Journal of Finance* 59.3, pp. 999–1037.
- Duarte, Victor et al. (Dec. 2021). *Simple Allocation Rules and Optimal Portfolio Choice Over the Lifecycle*. Working Paper 29559. National Bureau of Economic Research. DOI: [10.3386/w29559](https://doi.org/10.3386/w29559). URL: <http://www.nber.org/papers/w29559>.
- Duffie, Darrell et al. (1994). “Stationary Markov Equilibria”. In: *Econometrica: Journal of the Econometric Society*, pp. 745–781.
- Fernández-Villaverde, Jesús (Sept. 2025). *Deep Learning for Solving Economic Models*. Working Paper 34250. National Bureau of Economic Research. DOI: [10.3386/w34250](https://doi.org/10.3386/w34250). URL: <http://www.nber.org/papers/w34250>.

- Fu, Chao, Hsuan-Chih (Luke) Lin, and Atsuko Tanaka (Sept. 2025). *College Loans and Human Capital Investment*. Working Paper 34221. National Bureau of Economic Research. DOI: [10.3386/w34221](https://doi.org/10.3386/w34221). URL: <http://www.nber.org/papers/w34221>.
- Gârleanu, Nicolae and Stavros Panageas (2015). “Young, old, conservative, and bold: The implications of heterogeneity and finite lives for asset pricing”. In: *journal of political economy* 123.3, pp. 670–685.
- Geanakoplos, John, Michael Magill, and Martine Quinzii (2004). “Demography and the long-run predictability of the stock market”. In: *Brookings Papers on Economic Activity* 2004.1, pp. 241–325.
- Glover, Andrew et al. (2020). “Intergenerational Redistribution in the Great Recession”. In: *Journal of Political Economy* 128.10, pp. 3730–3778. DOI: [10.1086/708820](https://doi.org/10.1086/708820). URL: <https://doi.org/10.1086/708820>.
- Gomes, Francisco and Alexander Michaelides (2003). “Aggregate implications of defined benefit and defined contribution systems”. In: *Boston College Center for Retirement Research Working Paper* 2003-16.
- Gomes, Francisco, Alexander Michaelides, and Valery Polkovnichenko (2009). “Optimal savings with taxable and tax-deferred accounts”. In: *Review of Economic Dynamics* 12.4, pp. 718–735.
- Gopalakrishna, Goutham (2021). “Aliens and continuous time economies”. In: *Swiss Finance Institute Research Paper* 21-34.
- Gürkaynak, Refet S, Brian Sack, and Jonathan H Wright (2007). “The US Treasury yield curve: 1961 to the present”. In: *Journal of monetary Economics* 54.8, pp. 2291–2304.
- Han, Jiequn, Arnulf Jentzen, and Weinan E (2018). “Solving high-dimensional partial differential equations using deep learning”. In: *Proceedings of the National Academy of Sciences* 115.34, pp. 8505–8510.
- Han, Jiequn, Yucheng Yang, and Weinan E (2021). “DeepHAM: A Global Solution Method for Heterogeneous Agent Models with Aggregate Shocks”. In: DOI: [10.48550/ARXIV.2112.14377](https://arxiv.org/abs/2112.14377). URL: <https://arxiv.org/abs/2112.14377>.
- Hansen, Lars Peter and Ravi Jagannathan (1991). “Implications of security market data for models of dynamic economies”. In: *Journal of political economy* 99.2, pp. 225–262.

- Hasanhodzic, Jasmina (2015). “Borrowing Costs and the Equity Premium in Standard OLG Models”. In.
- Hayashi, Fumio (1982). “Tobin’s marginal  $q$  and average  $q$ : A neoclassical interpretation”. In: *Econometrica: Journal of the Econometric Society*, pp. 213–224.
- Heath, David, Robert Jarrow, and Andrew Morton (1992). “Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation”. In: *Econometrica: Journal of the Econometric Society*, pp. 77–105.
- Henriksen, Espen and Stephen Spear (2012). “Endogenous market incompleteness without market frictions: Dynamic suboptimality of competitive equilibrium in multiperiod overlapping generations economies”. In: *Journal of Economic Theory* 147.2, pp. 426–449.
- Hosseini, Roozbeh and Ali Shourideh (2019). “Retirement financing: An optimal reform approach”. In: *Econometrica* 87.4, pp. 1205–1265.
- Hull, John and Alan White (1990). “Pricing interest-rate-derivative securities”. In: *The review of financial studies* 3.4, pp. 573–592.
- Ionescu, Felicia (2009). “The federal student loan program: Quantitative implications for college enrollment and default rates”. In: *Review of Economic dynamics* 12.1, pp. 205–231.
- Jorgenson, Dale W (1963). “Capital theory and investment behavior”. In: *The American economic review* 53.2, pp. 247–259.
- Kim, Eungsik (2018). *Preference Heterogeneity, Aggregate Risk and the Welfare Effects of Social Security*. Tech. rep.
- Kingma, Diederik P and Jimmy Ba (2014). “Adam: A method for stochastic optimization”. In: *arXiv preprint arXiv:1412.6980*.
- Krueger, Dirk and Felix Kubler (2002). “Intergenerational risk-sharing via social security when financial markets are incomplete”. In: *American Economic Review* 92.2, pp. 407–410.
- (June 2006). “Pareto-Improving Social Security Reform when Financial Markets are Incomplete!?” In: *American Economic Review* 96.3, pp. 737–755. DOI: [10.1257/aer.96.3.737](https://doi.org/10.1257/aer.96.3.737). URL: <https://www.aeaweb.org/articles?id=10.1257/aer.96.3.737>.

- Kruse, Douglas L (1995). “Pension substitution in the 1980s: Why the shift toward defined contribution?”. In: *Industrial Relations: A Journal of Economy and Society* 34.2, pp. 218–241.
- Krusell, Per and Anthony A Smith Jr (1998). “Income and wealth heterogeneity in the macroeconomy”. In: *Journal of political Economy* 106.5, pp. 867–896.
- Laibson, David (1996). *Hyperbolic discount functions, undersaving, and savings policy*.
- Lochner, Lance and Alexander Monge-Naranjo (2011a). “Credit constraints in education”. In: *National Bureau of Economic Research*.
- (2011b). “The nature of credit constraints and human capital”. In: *American economic review* 101.6, pp. 2487–2529.
- Ma, Jennifer and Matea Pender (2021). *Trends in College Pricing and Student Aid 2021*. College Board.
- Maliar, Lilia, Serguei Maliar, and Pablo Winant (2021). “Deep learning for solving dynamic economic models.”. In: *Journal of Monetary Economics* 122, pp. 76–101.
- Merton, Robert C et al. (1971). “Theory of rational option pricing”. In.
- Morazzoni, Marta (2022). “Student Debt and Entrepreneurship in the US”. In.
- Moretto, Mauro (2021). “Essays on Overlapping Generations Models and Social Security”. PhD thesis. Carnegie Mellon University.
- Musiela, Marek and Marek Rutkowski (2005). *Martingale Methods in Financial Modelling*. 2nd ed. Vol. 36. Stochastic Modelling and Applied Probability. Berlin, Heidelberg: Springer. ISBN: 978-3-540-20966-9. DOI: [10.1007/b137866](https://doi.org/10.1007/b137866). URL: <https://link.springer.com/book/10.1007/b137866>.
- Raissi, Maziar, Paris Perdikaris, and George E Karniadakis (2019). “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations”. In: *Journal of Computational physics* 378, pp. 686–707.
- Reuter, Jonathan (2024). *Plan design and participant behavior in defined contribution retirement plans: Past, present, and future*. Tech. rep. National Bureau of Economic Research.
- Samuelson, Paul A (1958). “An exact consumption-loan model of interest with or without the social contrivance of money”. In: *Journal of Political Economy* 66.6, pp. 467–482.

- Shreve, Steven E (2004). *Stochastic calculus for finance II: Continuous-time models*. Vol. 11. Springer.
- Sirignano, Justin and Konstantinos Spiliopoulos (2018). “DGM: A deep learning algorithm for solving partial differential equations”. In: *Journal of computational physics* 375, pp. 1339–1364.
- Spear, Stephen (1988). “Existence and local uniqueness of functional rational expectations equilibria in dynamic economic models”. In: *Journal of Economic Theory* 44.1, pp. 124–155.
- Spear, Stephen and Sanjay Srivastava (1986). “Markov rational expectations equilibria in an overlapping generations model”. In: *Journal of Economic Theory* 38.1, pp. 35–62.
- Storesletten, Kjetil, Christopher I Telmer, and Amir Yaron (2007). “Asset pricing with idiosyncratic risk and overlapping generations”. In: *Review of Economic Dynamics* 10.4, pp. 519–548.
- Svensson, Lars E.O. (Sept. 1994). *Estimating and Interpreting Forward Interest Rates: Sweden 1992 - 1994*. Working Paper 4871. National Bureau of Economic Research. DOI: [10.3386/w4871](https://doi.org/10.3386/w4871). URL: <http://www.nber.org/papers/w4871>.
- Tobin, James (1969). “A general equilibrium approach to monetary theory”. In: *Journal of money, credit and banking* 1.1, pp. 15–29.
- Vasicek, Oldrich (1977). “An equilibrium characterization of the term structure”. In: *Journal of financial economics* 5.2, pp. 177–188.
- Yannelis, Constantine and Greg Tracey (2022). “Student loans and borrower outcomes”. In: *Annual Review of Financial Economics* 14, pp. 167–186.
- Yuana, Jun (2024). “Hedging Barrier Options Using Reinforcement Learning”. In: *Journal Of Investment Management* 22.4, pp. 16–25.
- Zhang, Adam (2022). *Before and After Target Date Investing: The General Equilibrium Implications of Retirement Savings Dynamics*. Tech. rep. Working Paper.